edudrift

An LLM's living hypothesis log on AI's impact on higher education

Santosh B. Srinivas · HEC Paris

Snapshot: 2026-05-13

https://santoshbs.github.io/theaidrift/edu/

edudrift is an experiment in letting a local open-weight large language model continuously process AI-related news signals and hypothesise their implications for specific higher-education capabilities. Each entry in this document is a working manuscript: it is generated and revised **entirely by the model**, with **no human in the editing loop**. The author's contribution is limited to designing the system — selecting RSS sources, defining the capability scaffold, writing the prompts, and committing the model's output to version control. The prose, the reasoning, the framing, and the choice of which signals matter are all the model's.

The capability scaffold is drawn from the **Higher Education Reference Models (HERM)** v3.2.0, a global industry-standard enterprise-architecture framework curated by the HERM Working Group of CAUDIT (the Council of Australasian University Directors of Information Technology), in collaboration with EDUCAUSE, UCISA, and EUNIS, and used by more than a thousand institutions worldwide. HERM is published under CC BY-NC-SA 4.0 and is available at https://www.caudit.edu.au/communities/caudit-higher-education-reference-models/.

For each capability, the log states three horizon hypotheses, each grounded in cited public signals:

- **Near** (2026–2031) — what current evidence most plausibly supports;
- **Mid** (2036–2041) — extrapolations contingent on signal trajectories holding;
- **Far** (2046–2051) — competing hypotheses where the trajectory is genuinely uncertain.

Hypotheses are abductive and revisable; each is version-controlled at the source repository. *Italic statements* are the formal hypothesis claims.

The live, continuously revised version is published at /theaidrift/edu/.

{*Caveat lector.* Because this document is fully model-generated, claims should be read as structured speculation grounded in cited signals, not as expert assessment. Citations point to the public sources the model drew on; readers should verify any claim against the original source before relying on it.}

Contents

Academic Administration **6**
Near (2026–2031) . . . 6
Mid (2036–2041) . . . 6
Far (2046–2051) . . . 6
References . . . 6

Business Capability Management **8**
Near (2026–2031) . . . 8
Mid (2036–2041) . . . 8
Far (2046–2051) . . . 8
References . . . 8

Commercial Delivery **10**
Near (2026–2031) . . . 10
Mid (2036–2041) . . . 10
Far (2046–2051) . . . 10
References . . . 10

Commercial Sourcing **11**
Near (2026–2031) . . . 11
Mid (2036–2041) . . . 11
Far (2046–2051) . . . 11
References . . . 11

Corporate Governance **13**
Near (2026–2031) . . . 13
Mid (2036–2041) . . . 13
Far (2046–2051) . . . 13
References . . . 14

Curriculum Management **16**
Near (2026–2031) . . . 16
Mid (2036–2041) . . . 16
Far (2046–2051) . . . 16
References . . . 17

Facilities & Property Management **19**
Near (2026–2031) . . . 19
Mid (2036–2041) . . . 19
Far (2046–2051) . . . 19
References . . . 20

Finance Management **21**
Near (2026–2031) . . . 21
Mid (2036–2041) . . . 21
Far (2046–2051) . . . 21
References . . . 22

Government, Public & Stakeholder Relationships **23**
Near (2026–2031) . . . 23
Mid (2036–2041) . . . 23
Far (2046–2051) . . . 23
References . . . 23

Human Resource Management **24**

Near (2026–2031) . . . 24
Mid (2036–2041) . . . 24
Far (2046–2051) . . . 24
References . . . 25

Information & Communication Technology Mgt **26**
Near (2026–2031) . . . 26
Mid (2036–2041) . . . 26
Far (2046–2051) . . . 26
References . . . 27

Information Management **30**
Near (2026–2031) . . . 30
Mid (2036–2041) . . . 30
Far (2046–2051) . . . 31
References . . . 31

Legal Services **37**
Near (2026–2031) . . . 37
Mid (2036–2041) . . . 37
Far (2046–2051) . . . 37
References . . . 38

Library Management **40**
Near (2026–2031) . . . 40
Mid (2036–2041) . . . 40
Far (2046–2051) . . . 40
References . . . 40

Promotions Management **41**
Near (2026–2031) . . . 41
Mid (2036–2041) . . . 41
Far (2046–2051) . . . 41
References . . . 42

Research Administration **43**
Near (2026–2031) . . . 43
Mid (2036–2041) . . . 43
Far (2046–2051) . . . 43
References . . . 44

Research Delivery **46**
Near (2026–2031) . . . 46
Mid (2036–2041) . . . 46
Far (2046–2051) . . . 46
References . . . 47

Research Funding **49**
Near (2026–2031) . . . 49
Mid (2036–2041) . . . 49
Far (2046–2051) . . . 49
References . . . 49

Research Impact **51**
Near (2026–2031) . . . 51
Mid (2036–2041) . . . 51
Far (2046–2051) . . . 51

References . . . 52
Research Improvement **53**
Near (2026–2031) . . . 53
Mid (2036–2041) . . . 53
Far (2046–2051) . . . 53
References . . . 54
Research Opportunity Analysis **55**
Near (2026–2031) . . . 55
Mid (2036–2041) . . . 55
Far (2046–2051) . . . 55
References . . . 55
Research Publications **56**
Near (2026–2031) . . . 56
Mid (2036–2041) . . . 56
Far (2046–2051) . . . 56
References . . . 57
Research Training **58**
Near (2026–2031) . . . 58
Mid (2036–2041) . . . 58
Far (2046–2051) . . . 58
References . . . 59
Strategy Management **60**
Near (2026–2031) . . . 60
Mid (2036–2041) . . . 60
Far (2046–2051) . . . 60
References . . . 61
Student Administration **62**
Near (2026–2031) . . . 62
Mid (2036–2041) . . . 62
Far (2046–2051) . . . 62
References . . . 62
Student Admission Management **63**
Near (2026–2031) . . . 63
Mid (2036–2041) . . . 63
Far (2046–2051) . . . 63
References . . . 63
Student Assessment **64**
Near (2026–2031) . . . 64
Mid (2036–2041) . . . 64
Far (2046–2051) . . . 64
References . . . 65
Student Attraction & Recruitment **67**
Near (2026–2031) . . . 67
Mid (2036–2041) . . . 67
Far (2046–2051) . . . 67
References . . . 68
Student Enrolment **69**

Near (2026–2031) . . . 69
Mid (2036–2041) . . . 69
Far (2046–2051) . . . 69
References . . . 69
Student Support & Wellbeing Management **70**
Near (2026–2031) . . . 70
Mid (2036–2041) . . . 70
Far (2046–2051) . . . 70
References . . . 71
Supporting Services **72**
Near (2026–2031) . . . 72
Mid (2036–2041) . . . 72
Far (2046–2051) . . . 72
References . . . 72
Teaching & Learning Delivery **74**
Near (2026–2031) . . . 74
Mid (2036–2041) . . . 74
Far (2046–2051) . . . 74
References . . . 75

Academic Administration

capability_id: academic-administration · created: 2026-04-10 · revised: 2026-04-10 · refs: 2

Near (2026–2031)

For Academic Administration broadly, evidence that AI-enabled misconduct overlaps with traditional forms of plagiarism [1] may drive a shift in academic policy & regulation management toward integrated behavioral frameworks rather than tool-specific bans [1]. For institutions implementing structured faculty development partnerships, as demonstrated at the University of Central Florida [2], this transition may be supported by bottom-up infrastructure involving collaborations between teaching centers and digital learning divisions [2]. This represents a change in educational substance by refocusing regulation on behavioral patterns of dishonesty rather than the specific technology employed. Because updates to academic policy typically require faculty senate approval and governance review, the pace of institutional adoption may remain slow. Processes such as academic year scheduling and timetabling management are unlikely to change in this period as they are not functionally linked to academic integrity signals.

Hypothesis: In the near term, academic policy & regulation management may shift from tool-specific AI bans toward holistic integrity frameworks that treat AI use as one of several overlapping plagiarism behaviors.

Mid (2036–2041)

For Academic Administration broadly, if current trajectories hold, the identified link between AI use and general plagiarism [1] could lead to the development of longitudinal behavioral profiles within academic policy & regulation management. If institutions continue to build evolving policy infrastructures through cross-functional partnerships [2], regulatory mechanisms might shift from evaluating single instances of misconduct to identifying patterns of dishonesty across a student's academic career. This transition would likely be constrained by data privacy laws and institutional due process requirements, which tend to persist regardless of technological shifts. The fundamental administrative structure of the academic year remains unlikely to change, as it is governed by external labor contracts and seasonal student demographics rather than integrity regulation.

Hypothesis: In the mid term, academic policy & regulation management may move toward a behavioral-pattern model of integrity enforcement, extrapolating from current evidence of overlapping plagiarism methods.

Far (2046–2051)

For Academic Administration broadly, the long-term trajectory of integrity regulation remains uncertain, with two primary rival outcomes. One possibility is that the persistence of overlapping plagiarism [1] leads to the automation of the disciplinary process through algorithmic pattern matching. A competing outcome is that the continued erosion of traditional authorship leads to a collapse of the current regulatory regime, where the "product" of education is no longer a verifiable independent work. It remains unknown how professional accreditation bodies will redefine "competence" if traditional plagiarism regulations become obsolete. Even at this horizon, the legal requirement for human-led due process in high-stakes disciplinary appeals is likely to resist full automation.

Hypothesis A: In the far term, academic policy & regulation management may evolve into a behavioral monitoring system that identifies integrity risks.

Hypothesis B (competing): In the far term, the capability of academic policy & regulation management may be fundamentally diminished as the institution shifts away from the validation of independent student authorship.

References

1. The Conversation Au (2026-04-01). "Does AI mean more uni students are plagiarising their work?" https://theconversation.com/does-ai-mean-more-uni-students-are-plagiarising-their-work-279565

2. EDUCAUSE Review (2026-04-08). “Building AI Initiatives with Bottom-Up Champions” https://er.educause.edu/articles/2026/4/building-ai-initiatives-with-bottom-up-champions

Business Capability Management

capability_id: business-capability-management · created: 2026-05-06 · revised: 2026-05-06 · refs: 5

Near (2026–2031)

For business capability management broadly, the adoption of AI operating models requires a redesign of organisational design and change portfolio management to prevent governance sprawl as institutions functionally transition toward software-factory operations [4]. For a subsegment of institutions deploying AI infrastructure, a low success rate—where only 28% of projects fully realize return on investment—may force benefits management to shift in substance toward more stringent ROI validation [3]. The delivery mechanism of programme & project management is altered by the adoption of specialized AI lifecycle frameworks designed to bridge the gap between conceptual prototypes and production readiness [1]. In institutions utilizing Scaled Agile Frameworks (SAFe), the delivery mechanism of programme & project management is further modified through the automation of delivery metrics [5]. In institutions focusing on rapid tool deployment, change portfolio management may shift in substance toward human-centric adoption strategies to avoid the cost of unused workflows [2]. Quality management processes are unlikely to change yet, as they rely on stable administrative specifications that remain independent of AI implementation.

Hypothesis: In the near term, business capability management will integrate AI-specific operating models and lifecycle frameworks to govern the transition toward software-factory-style operations and mitigate high infrastructure failure rates.

Mid (2036–2041)

For business capability management broadly, if the trend of low ROI for infrastructure persists, the capability may shift from experimental AI adoption to disciplined architectural integration [3]. Extrapolating from the necessity of governing AI sprawl [4], enterprise architecture could evolve to standardize the software-factory model across the institution. This represents a change in the educational substance of the capability, as institutional value is measured by production-level operational integration rather than prototype capability [1]. The fundamental mechanisms of institutional accountability—determining legal and professional responsibility for capability failure—will likely persist unchanged due to regulatory and governance constraints.

Hypothesis: In the mid term, the capability may shift toward a disciplined architectural integration where AI-driven operating models are standardized as core institutional utilities.

Far (2046–2051)

For business capability management broadly, the long-term trajectory remains uncertain regarding the automation of strategic decision-making. One outcome is that change portfolio management becomes largely automated, with AI systems optimizing resource allocation based on real-time adoption and ROI data [2][3]. However, the process of determining the social legitimacy of institutional goals resists change even at this horizon and likely remains a human-centric function. It remains unknown whether AI can manage benefits management without replicating the systemic ROI failures observed in early infrastructure projects [3].

Hypothesis A: In the far term, business capability management may become largely automated, with AI systems optimizing the change portfolio through real-time performance feedback.

Hypothesis B (competing): In the far term, the capability may return to highly centralized human control as a defensive measure against the systemic unpredictability of autonomous governance models.

References

1. Dzone Ai (2026-04-03). "From Concept to Production: A Strategic Framework for AI/ML Project Success." https://feeds.dzone.com/link/23558/17313177/ai-lifecycle-framework-for-successful-projects

2. Dev To Ai (2026-04-06). "How to Build an AI Workflow That Your Team Will Actually Use." https://dev.to/maxothex/how-to-build-an-ai-workflow-that-your-team-will-actually-use-55he
3. The Register (2026-04-07). "Only 28% of AI infrastructure projects fully pay off, survey finds." https://go.theregister.com/feed/www.theregister.com/2026/04/07/ai_returns_gartner/
4. Dev To Ai (2026-04-11). "AI Operating Models: The $2M Workflow Redesign Your Board Isn't Asking About." https://dev.to/dr_hernani_costa/ai-operating-models-the-2m-workflow-redesign-your-board-isnt-asking-about-1dfd
5. Dzone Ai (2026-04-22). "Revolutionizing Scaled Agile Frameworks with AI, MuleSoft, and AWS: An Insider's Perspective." https://feeds.dzone.com/link/23558/17323804/revolutionizing-scaled-agile-frameworks-with-ai

Commercial Delivery

capability_id: commercial-delivery · created: 2026-04-09 · revised: 2026-04-09 · refs: 1

Near (2026–2031)

For US-based institutions utilizing blacklisted AI technology, the process of commercial activity exit management may be triggered by federal regulatory actions and judicial rulings [1]. This represents a change in the delivery mechanism of risk management, shifting exit protocols from planned lifecycle sunsets to forced compliance measures. Because the current signal is tied to specific litigation, other processes such as commercial activity output management are unlikely to change in the near term. The mechanism driving this shift is the legal necessity to terminate commercial ties with entities designated as prohibited by the state [1].

Hypothesis: In the near term, a subset of US institutions will experience an increase in forced commercial activity exit management events due to geopolitical trade restrictions on specific AI vendors.

Mid (2036–2041)

For US-based institutions, if current trajectories of geopolitical technology blacklisting hold, the process of commercial activity exit management may become a standardized component of commercial risk auditing. Extrapolating from recent litigation [1], institutions could implement more frequent "compliance triggers" within their commercial activity issue management processes to avoid abrupt service terminations. The educational substance of the commercial activities remains unchanged, but the administrative delivery of vendor oversight becomes more conservative. The fundamental requirement for legal counsel in commercial delivery processes will persist unchanged due to the complexity of trade law.

Hypothesis: In the mid term, US institutions may adopt a more modular approach to commercial activity exit management to reduce the institutional shock of regulatory blacklisting.

Far (2046–2051)

For US-based institutions, the long-term state of commercial activity exit management depends on the stability of international trade regimes. One plausible outcome is the emergence of fragmented "technological blocs," where commercial delivery processes are strictly bifurcated by national security boundaries, making the exit management of foreign-developed AI a routine operational task. A competing outcome is a return to globalized standards where such blacklists are rendered obsolete by new international treaties. It remains unknown whether future regulatory frameworks will provide "grace periods" for institutional exits or mandate immediate cessation. The basic institutional need to manage the legal termination of contracts will resist change regardless of the technology involved.

Hypothesis A: In the far term, commercial activity exit management becomes a highly automated, trigger-based process within a fragmented global AI market.

Hypothesis B (competing): In the far term, the volatility of blacklisting ceases as global commercial delivery standards are re-harmonized, reducing the frequency of forced exit management events.

References

1. Ars Technica (2026-04-09). "Trump-appointed judges refuse to block Trump blacklisting of Anthropic AI tech." https://arstechnica.com/tech-policy/2026/04/trump-appointed-judges-refuse-to-block-trump-blacklisting-of-anthropic-ai-tech/

Commercial Sourcing

capability_id: commercial-sourcing · created: 2026-04-09 · revised: 2026-04-09 · refs: 2

Near (2026–2031)

For a subset of institutions using AI-driven search tools and those operating within US regulatory jurisdictions, the reliability of commercial market analysis may be compromised by both external manipulation and geopolitical constraints. AI-focused SEO tactics attempt to influence LLM-generated lists of vendors, which may introduce bias into the process of commercial opportunity identification [1]. Simultaneously, the blacklisting of specific high-capability AI providers, such as Anthropic, restricts the pool of viable vendors available for commercial opportunity assessment in the US [2]. This represents a change to the delivery mechanism of vendor discovery via AI SEO and a regulatory constraint on the available market via blacklisting. Because procurement processes often rely on initial discovery lists to narrow the field, these factors could lead to the exclusion of non-optimized or legally restricted vendors [1][2]. However, formal procurement mandates requiring a predefined set of competitive bids are unlikely to change, as these regulatory constraints prevent total reliance on AI-generated suggestions.

Hypothesis: In the near term, for the affected subsegments and regions, commercial market analysis may become less reliable due to a combination of manipulated AI responses and regulatory restrictions on vendor availability.

Mid (2036–2041)

For these same subsets and regions, if current trajectories of AI-driven SEO and geopolitical blacklisting hold, the process of commercial opportunity assessment could shift toward a more adversarial verification model. Institutions may introduce mandatory cross-referencing of AI-generated vendor lists against curated, compliance-verified industry databases to mitigate the risk of both manipulated rankings and the selection of prohibited entities [1][2]. This would introduce a formal validation step into the commercial market analysis process to ensure the integrity and legality of the source data. Despite this, the internal institutional governance regarding the final selection of vendors is likely to persist unchanged due to high financial stakes and existing audit requirements.

Hypothesis: In the mid term, the commercial market analysis process for these subsegments may integrate formal validation and compliance steps to counter misinformation and regulatory risk.

Far (2046–2051)

For institutions relying on AI for sourcing, the outcome depends on whether AI search engines develop verifiable transparency standards and whether geopolitical vendor restrictions stabilize. One plausible outcome is that commercial opportunity identification becomes a specialized technical function where AI-driven discovery is viewed as an unreliable lead-generation step rather than a valid analysis tool. It remains unknown whether AI SEO will evolve into a standardized digital signaling mechanism or if blacklisting becomes a permanent feature of institutional procurement. The underlying need for human-led final due diligence will likely resist change because of the legal liability associated with institutional commercial contracts.

Hypothesis A: In the far term, commercial market analysis for this subsegment may decouple from generative AI search, reverting to trusted, closed-loop industry registries to avoid noise and regulatory volatility.

Hypothesis B (competing): In the far term, a new standard for verifiable AI sourcing and automated compliance filtering may emerge, allowing institutions to efficiently identify viable vendors while filtering out manipulated content.

References

1. The Verge (2026-04-06). "Can AI responses be influenced? The SEO industry is trying." https://www.theverge.com/tech/900302/ai-seo-industry-google-search-chatgpt-gemini-marketing

2. Ars Technica (2026-04-09). "Trump-appointed judges refuse to block Trump blacklisting of Anthropic AI tech." https://arstechnica.com/tech-policy/2026/04/trump-appointed-judges-refuse-to-block-trump-blacklisting-of-anthropic-ai-tech/

Corporate Governance

capability_id: corporate-governance · created: 2026-05-13 · revised: 2026-05-13 · refs: 23

Near (2026–2031)

For Corporate Governance broadly, the transition from generative AI to autonomous agent execution shifts the educational substance of risk management from the validation of model outputs to the governance of executed actions [5]. To mitigate the risk of autonomous failures, risk management processes are integrating sandbox execution environments [24] and runtime security toolkits [10] to constrain agent autonomy. Policy and regulation management is adopting agent registries [13] and autonomous governance tools [2] to detect and regulate "shadow AI" deployments. Within engineering and security subsegments, risk management is incorporating frameworks for AI-specific risk tiering and asset inventory [32]. In the European Union, compliance monitoring and reporting is constrained by mandates to produce Annex IV technical files [8], justify data processing to the European Data Protection Board [17], and manage agentic AI under the EU AI Act [12]. In the United States, policy and regulation management is adapting to a proposed national AI legislative framework [26] while managing risks associated with executive-led technology blacklisting [14] and criminal liability probes following agent-linked incidents [15][28][29]. Within specialized subsegments, risk management processes are incorporating "judgment layers" for financial agents to prevent automated outputs from being treated as final decisions [19], as well as guards against data bias in hiring [3] and "LLM nepotism" in organizational evaluation [20]. The fundamental legal principle that an institution cannot delegate fiduciary liability for a failure to a software agent is unlikely to change [1].

Hypothesis: In the near term, institutions will shift risk and policy frameworks to mitigate the liability of authorized agents and the operational risks of shadow AI, while moving compliance monitoring toward runtime enforcement and regulatory technical documentation.

Mid (2036–2041)

For Corporate Governance broadly, if current trajectories of agent autonomy and regulatory tension hold, internal audit and reporting may evolve from retrospective periodic sampling to the use of governance toolkits that provide immutable, real-time audit trails [6]. This mechanism could potentially reduce the attribution problem that currently complicates the assignment of legal responsibility for autonomous agent failures [1]. Extrapolating from the emergence of AI operating models, the substance of board-level oversight may shift toward the continuous redesign of high-cost organizational workflows rather than the approval of static policies [18]. In regions where strict regulatory regimes persist, compliance monitoring will likely remain a high-friction process requiring manual human certification despite automated data collection [12]. The requirement for human-in-the-loop verification for high-stakes institutional decisions is unlikely to change due to social legitimacy constraints.

Hypothesis: In the mid term, the capability may shift from periodic auditing to real-time governance monitoring, with board oversight focusing on the architectural design of AI-integrated workflows.

Far (2046–2051)

For Corporate Governance broadly, the interaction between autonomous agents and business continuity management remains a primary uncertainty due to the persistence of systems that cannot explain their own decision-making logic [25]. This uncertainty is compounded by evidence that "secret loyalties" can be trained into models to bypass standard black-box audits, potentially undermining the reliability of automated compliance monitoring [30]. One plausible outcome is the institutionalization of a "seatbelt" approach, where governance is an automated, preventative layer that overrides agent action based on hard-coded constraints [4]. A competing outcome involves a move toward total transparency mandates, though it remains unknown if technical solutions for full explainability will ever emerge to satisfy regulatory requirements. The basic necessity for a defined legal person or entity to hold ultimate accountability for institutional actions will likely resist change regardless of the level of agent autonomy.

Hypothesis A: In the far term, corporate governance may rely on "seatbelt" architectures that use hard-coded constraints to override autonomous agent actions.

Hypothesis B (competing): In the far term, governance may fragment as a result of the inability to audit "black box" agent loyalties, leading to a return to highly restrictive, human-centric control regimes.

References

1. The Register (2026-04-05). "AI agents promise to 'run the business,' but who is liable if things go wrong?" https://go.theregister.com/feed/www.theregister.com/2026/04/05/ai_agents_liability/
2. Ai News (2026-04-02). "KiloClaw targets shadow AI with autonomous agent governance" https://www.artificialintelligence-news.com/news/kiloclaw-targets-shadow-ai-autonomous-agent-governance/
3. Ai Trends (2021-10-28). "Promise and Perils of Using AI for Hiring: Guard Against Data Bias" https://www.aitrends.com/ai-world-government/promise-and-perils-of-using-ai-for-hiring-guard-against-data-bias/
4. Dev To Ai (2026-04-06). "AI Doesn't Need a Bigger Engine. It Needs a Seatbelt." https://dev.to/tomleelive/ai-doesnt-need-a-bigger-engine-it-needs-a-seatbelt-5k8
5. Ai News (2026-04-06). "As AI agents take on more tasks, governance becomes a priority" https://www.artificialintelligence-news.com/news/as-ai-agents-take-on-more-tasks-governance-becomes-a-priority/
6. Dev To Ai (2026-04-06). "AI agent governance tools compared - 2026 landscape" https://dev.to/jagmarques/ai-agent-governance-tools-compared-2026-landscape-53hm
7. Dev To Ai (2026-04-07). "What's actually inside an EU AI Act Annex IV technical file (and what engineers usually miss)" https://dev.to/rostyslav_myronenko_faf2f/whats-actually-inside-an-eu-ai-act-annex-iv-technical-file-and-what-engineers-usually-miss-4nh4
8. Ai News (2026-04-08). "Microsoft open-source toolkit secures AI agents at runtime" https://www.artificialintelligence-news.com/news/microsoft-open-source-toolkit-secures-ai-agents-at-runtime/
9. Ai News (2026-04-09). "Agentic AI's governance challenges under the EU AI Act in 2026" https://www.artificialintelligence-news.com/news/agentic-ais-governance-challenges-under-the-eu-ai-act-in-2026/
10. The Register (2026-04-09). "AWS: Agents shouldn't be secret, so we built a registry for them" https://go.theregister.com/feed/www.theregister.com/2026/04/09/aws_ai_agent_registry/
11. Ars Technica (2026-04-09). "Trump-appointed judges refuse to block Trump blacklisting of Anthropic AI tech" https://arstechnica.com/tech-policy/2026/04/09/trump-appointed-judges-refuse-to-block-trump-blacklisting-of-anthropic-ai-tech/
12. TechCrunch (2026-04-09). "Florida AG to probe OpenAI, alleging possible connection to FSU shooting" https://techcrunch.com/2026/04/09/florida-ag-to-probe-openai-alleging-possible-connection-to-fsu-shooting/
13. Dev To Ai (2026-04-10). "The EDPB Is Asking About Your AI Agents. Most Teams Can't Answer." https://dev.to/waxell/the-edpb-is-asking-about-your-ai-agents-most-teams-cant-answer-gfk
14. Dev To Ai (2026-04-11). "AI Operating Models: The $2M Workflow Redesign Your Board Isn't Asking About" https://dev.to/dr_hernani_costa/ai-operating-models-the-2m-workflow-redesign-your-board-isnt-asking-about-1dfd
15. Dev To Ai (2026-04-11). "Judgment Layer for Financial AI Agents" https://dev.to/thatechmaestro/judgment-layer-for-financial-ai-agents-141
16. Arxiv Cs Cy (2026-04-14). "LLM Nepotism in Organizational Governance" https://arxiv.org/abs/2604.09620
17. Ai News (2026-04-16). "OpenAI Agents SDK improves governance with sandbox execution" https://www.artificialintelligence-news.com/news/openai-agents-sdk-improves-governance-sandbox-execution/
18. Hackernoon Ai (2026-04-21). "The Hidden Risk of AI Agents: Systems That Can't Explain Themselves" https://hackernoon.com/the-hidden-risk-of-ai-agents-systems-that-cant-explain-

themselves?source=rss
19. EDUCAUSE Review (2026-04-21). "White House Unveils National AI Legislative Framework" https://er.educause.edu/articles/2026/4/white-house-unveils-national-ai-legislative-framework
20. BBC News (2026-04-21). "OpenAI faces criminal probe over role of ChatGPT in shooting" https://www.bbc.com/news/articles/c62j4ldp2jqo?at_medium=RSS&at_campaign=rss
21. NPR (2026-04-21). "Florida AG launches criminal investigation into ChatGPT over FSU shooting" https://www.npr.org/2026/04/21/nx-s1-5793967/florida-openai-investigation-mass-shooting-fsu
22. Lesswrong (2026-04-22). "Narrow Secret Loyalty Dodges Black-Box Audits" https://www.lesswrong.com/posts/EzdgPbewjeTNHA5F3/narrow-secret-loyalty-dodges-black-box-audits
23. Marktechpost (2026-04-23). "Mend.io Releases AI Security Governance Framework Covering Asset Inventory, Risk Tiering, AI Supply Chain Security, and Maturity Model" https://www.marktechpost.com/2026/04/23/mend-io-releases-ai-security-governance-framework-covering-asset-inventory-risk-tiering-ai-supply-chain-security-and-maturity-model/

Curriculum Management

capability_id: curriculum-management · created: 2026-05-10 · revised: 2026-05-10 · refs: 17

Near (2026–2031)

For curriculum management broadly, the process of curriculum design is shifting toward the integration of workforce-aligned AI skillsets [12][14] and a change in educational substance toward skills-based frameworks [13] that prioritize judgment and process over the production of static academic artifacts [18]. In this broad scope, curriculum & resource development is adopting transparency frameworks to document the role of AI in the creation of educational content [5]. For a subset of institutions, specifically those in Career and Technical Education (CTE) and software engineering, the delivery mechanism of curriculum & resource development is shifting toward the use of automated multimodal content generation [6], localization [3], and structured AI integration into domain-specific curricula [15]. Within these subsegments, new processes for curriculum change management are emerging to govern the accuracy, quality, and trust of AI-generated content [10] and to maintain the validity of associated assessments [16]. For institutions implementing structured faculty development, as demonstrated at Indiana Wesleyan University [1] and through bottom-up champion networks [8], the process of professional learning (staff) is utilizing "course refresh" institutes to update educational substance [1][8], with a focus on building educator capacity rather than tool mastery [2]. In regions where teaching standards are being updated, such as the United States, curriculum design for educator preparation is being revised to incorporate updated professional standards [17]. For a subset of institutions, the delivery mechanism of professional learning (staff) is incorporating microlearning via LMS to improve the adoption of new standards [19]. In certain public institutions, as seen at Iowa State University [11], the process of curriculum retirement management is being triggered by mandated institutional reviews. Professional accreditation processes are unlikely to change significantly in this window because accreditation bodies typically operate on multi-year review cycles that lag behind classroom-level revisions.

Hypothesis: In the near term, curriculum management may shift toward integrating workforce-aligned skills and transparency frameworks in resource development, supported by a mix of structured faculty refreshes and microlearning for staff capacity.

Mid (2036–2041)

For the subset of institutions and subsegments identified in the near term, if current trajectories hold, curriculum change management may transition from ad-hoc updates to systemic, recurring reviews of learning experiences. Extrapolating from "course refresh" and bottom-up capacity models, the process of professional learning (staff) could become a permanent, cyclical requirement rather than a periodic intervention [1][2][8]. In CTE and technical contexts, delivery-mechanism efficiency gains in curriculum & resource development may increase the frequency of curriculum retirement management as specific technical skills obsolesce more rapidly [3][14]. However, the fundamental governance structures—such as faculty senate approvals and committee-based curriculum reviews—will likely persist unchanged to maintain institutional academic legitimacy.

Hypothesis: In the mid term, institutions with established development structures may formalize cyclical curriculum refreshes and accelerated resource updates as core institutional processes.

Far (2046–2051)

For these specific institutional contexts, the mechanism of curriculum design may diverge into rival operational modes. One outcome involves the emergence of highly modular curriculum design where individual learning experiences are updated in near real-time based on technical workforce signals, potentially eroding the stability of a cohesive degree programme. A rival outcome involves a return to highly stable, human-centric curriculum design that emphasizes timeless cognitive capabilities as a hedge against technical volatility [7]. It remains unknown whether professional accreditation bodies will adapt to real-time modularity or mandate the preservation of fixed curricula. The social legitimacy of the degree as a signal

of completed, standardized study resists change even at this horizon.

Hypothesis A: In the far term, some institutions may move toward hyper-modular, signal-driven curriculum design that updates in near real-time.

Hypothesis B (competing): In the far term, institutions may pivot toward "slow curriculum" models that prioritize enduring cognitive substance over rapidly shifting technical skills.

References

1. EDUCAUSE Review (2026-03-31). "Advancing AI Literacy: A Faculty Course Refresh Institute at Indiana Wesleyan University." https://er.educause.edu/articles/2026/3/advancing-ai-literacy-a-faculty-course-refresh-institute-at-indiana-wesleyan-university
2. Getting Smart (2026-03-27). "AI Literacy is Not Tool Mastery: How to Build Sustained Educator Capacity." https://www.gettingsmart.com/2026/03/27/ai-literacy-is-not-tool-mastery-how-to-build-sustained-educator-capacity/
3. Elearning Industry (2026-04-05). "Leveraging AI And Automation To Streamline CTE Content Creation And Localization." https://elearningindustry.com/leveraging-ai-and-automation-to-streamline-cte-content-creation-and-localization
4. Higher Ed Dive (2026-04-07). "Colleges ramp up offerings to teach students to be AI ethicists." https://www.highereddive.com/news/colleges-programs-ai-ethicists/816703/
5. Elearning Industry (2026-04-07). "A 4-Step Framework For Using AI Transparently In Educational Content." https://elearningindustry.com/a-4-step-framework-for-using-ai-transparently-in-educational-content
6. Arxiv Cs Cy (2026-04-08). "EduIllustrate: Towards Scalable Automated Generation Of Multimodal Educational Content." https://arxiv.org/abs/2604.05005
7. Dev To Ai (2026-04-08). "What Even Is The Purpose of Learning?" https://dev.to/mateebhussain/what-even-is-the-purpose-of-learning-4c2e
8. EDUCAUSE Review (2026-04-08). "Building AI Initiatives with Bottom-Up Champions." https://er.educause.edu/articles/2026/4/building-ai-initiatives-with-bottom-up-champions
9. Arxiv Cs Cy (2026-04-09). "Thinking in Graphs with CoMAP: A Shared Visual Workspace for Designing Project-Based Learning." https://arxiv.org/abs/2604.06200
10. Elearning Industry (2026-04-10). "AI-Generated Learning Content Governance: Ensuring Quality, Accuracy, And Trust." https://elearningindustry.com/ai-generated-learning-content-governance-ensuring-quality-accuracy-and-trust
11. Higher Ed Dive (2026-04-10). "Iowa State plans to cut or merge 23 programs after mandated review." https://www.highereddive.com/news/iowa-state-university-program-cuts-review-regents/817262/
12. Higher Ed Dive (2026-04-16). "AI is remaking the workforce. How can colleges ensure students thrive?" https://www.highereddive.com/news/ai-is-remaking-the-workforce-how-can-colleges-ensure-students-thrive/817598/
13. Elearning Industry (2026-04-15). "The Case For Skills-Based Learning In A Rapidly Changing Workplace." https://elearningindustry.com/the-case-for-skills-based-learning-in-a-rapidly-changing-workplace
14. Higher Ed Dive (2026-04-21). "Employers say they struggle to find graduates with the right AI skillset." https://www.highereddive.com/news/employers-struggle-to-find-workers-with-the-right-ai-skillset/818017/
15. Arxiv Cs Cy (2026-04-21). "Bringing AI into the Classroom: A Structured Approach for Integrating AI into Software Engineering Education." https://arxiv.org/abs/2604.16328
16. Elearning Industry (2026-04-22). "AI Assessment Guardrails: How To Use AI Without Breaking Validity And Trust." https://elearningindustry.com/ai-assessment-guardrails-how-to-use-ai-without-breaking-validity-and-trust
17. Getting Smart (2026-04-23). "America is Updating Teaching Standards Without a Clear Picture of the Future Educator." https://www.gettingsmart.com/2026/04/23/america-is-updating-teaching-

standards-without-a-clear-picture-of-the-future-educator/
18. EDUCAUSE Review (2026-04-23). "Teaching and Learning with AI: Three Priorities." https://er.educause.edu/podcasts/educause-exchange/teaching-and-learning-with-ai-three-priorities
19. Elearning Industry (2026-04-23). "Microlearning LMS: Does It Actually Improve Adoption?" https://elearningindustry.com/microlearning-lms-does-it-actually-improve-adoption
20. Elearning Industry (2026-04-23). "How AI Prompts Are Redefining Creativity In Instructional Design." https://elearningindustry.com/how-ai-prompts-are-redefining-creativity-in-instructional-design

Facilities & Property Management

capability_id: facilities-property-management · created: 2026-05-11 · revised: 2026-05-11 · refs: 8

Near (2026–2031)

For a subset of institutions adopting physical AI and for those in US regions subject to emerging land-use policies, the delivery mechanism for property maintenance, campus security, and estate development may shift [1][3][4]. In property maintenance, the adoption of robot mowers may replace specific manual labor tasks in groundskeeping [1], while AI-driven transcription for technician voice notes may streamline documentation workflows for facility repairs [2]. Sensing capabilities in robotic systems, such as reading gauges and thermometers, may automate the inspection of industrial plant equipment [5]. For campus security and space utilisation, the deployment of physical AI may alter perimeter monitoring [3], while AI video analytics may transform traditional CCTV into intelligent monitoring systems capable of detecting real-time risks and analyzing movement patterns [8]. Predictive intelligence in large venues may change how facility conditions are managed through the forecasting of behavioral patterns [7]. In estate development and management, AI platforms may accelerate the material takeoff process by quantifying materials for project budgeting [6], though state and local policy interventions in the US may constrain the siting and expansion of AI-specific data center facilities [4]. Processes such as mail management and commercial tenancy are unlikely to change yet because they remain dependent on manual administrative protocols and established contractual law.

Hypothesis: In the near term, a subset of institutions may reduce operational labor in maintenance and security through targeted automation while facing increased regional regulatory constraints on the development of AI-specific physical infrastructure.

Mid (2036–2041)

For institutions implementing these systems, if current trajectories in predictive intelligence and robotics hold, the scope of automated surveillance and maintenance may expand while estate development for AI facilities remains contingent on local policy compliance [3][4][8]. The process of property maintenance may shift toward continuous robotic auditing of facility health as sensing capabilities for industrial equipment mature [5]. In estate development, the use of AI for project estimation may move from material takeoffs to more integrated budgetary forecasting [6], while the development of data centers may require more complex negotiations with local municipalities regarding resource extraction [4]. Robotics in groundskeeping may move beyond mowing to encompass a broader range of repetitive exterior vegetation management [1]. Long-term campus planning and architectural synthesis are likely to remain human-led because these processes require alignment with institutional strategy and aesthetic values.

Hypothesis: In the mid term, routine exterior maintenance, industrial plant inspection, and perimeter security may become predominantly automated for early adopters, while the development of AI-centric facilities faces heightened regional regulatory barriers.

Far (2046–2051)

For the subsegments of exterior maintenance, industrial inspection, and estate development, outcomes diverge based on the tension between automation efficiency and environmental or social policy [1][4][5]. One scenario involves a transition to a highly automated facility management regime where human presence is limited to system oversight and both the physical perimeter and internal industrial plants are managed by integrated AI agents [3][5][8]. A competing scenario involves a shift toward sustainable "wild" landscaping and strict regulatory prohibitions on energy-intensive AI infrastructure, rendering both robotic mowing and expansive data center development obsolete in favor of low-impact land use [1][4]. The extent to which AI-driven security systems erode social trust or campus openness remains an unresolvable uncertainty at this horizon. Core legal frameworks governing property ownership and commercial tenancy are likely to resist change even at this horizon.

Hypothesis A: In the far term, early adopting institutions may transition to a regime of bounded autonomy for facility management, where AI agents oversee the majority of physical plant and security operations.

Hypothesis B (competing): In the far term, environmental mandates and regulatory pushback against AI energy consumption may force a reversal of automated infrastructure expansion, shifting the capability toward low-tech, sustainable estate management.

References

1. Wired (2026). "Robot Mowers Are Actually Good Now." https://www.wired.com/story/robot-mowers-are-good-now-rave/
2. Dev To Ai (2026). "From Mumbles to Memos: Teaching AI to Understand Technician Voice Notes." https://dev.to/ken_deng_ai/from-mumbles-to-memos-teaching-ai-to-understand-technician-voice-notes-ckj
3. Ai News (2026). "Asylon and Thrive Logic bring physical AI to enterprise perimeter security." https://www.artificialintelligence-news.com/news/physical-ai-security-at-the-enterprise-perimeter-takes-a-step-closer/
4. Ai Now Institute (2026). "North Star Data Center Policy Toolkit: State and Local Policy Interventions to Stop Rampant AI Data Center Expansion." https://ainowinstitute.org/publications/data-center-policy-guide
5. Ars Technica (2026). "Boston Dynamics' robot dog now reads gauges and thermometers with Google's AI." https://arstechnica.com/ai/2026/04/robot-dogs-now-read-gauges-and-thermometers-using-google-gemini/
6. Ai News (2026). "Bobyard 2.0 offers improved takeoffs and unified AI for estimators." https://www.artificialintelligence-news.com/news/bobyard-2-0-offers-improved-takeoffs-and-unified-ai-for-estimators/
7. Dev To Ai (2026). "Architecting Predictive Intelligence for Smart Venues." https://dev.to/tanisha_khiara_4b390b7946/architecting-predictive-intelligence-for-smart-venues-1lkj
8. Dev To Ai (2026). "Integrating AI Video Analytics into Venue Operations." https://dev.to/tanisha_khiara_4b390b7946/integrating-ai-video-analytics-into-venue-operations-2m09

Finance Management

capability_id: finance-management · created: 2026-05-08 · revised: 2026-05-08 · refs: 8

Near (2026–2031)

For Finance Management broadly, the execution of financial decisions within AI-integrated workflows may incorporate dedicated judgment layers to prevent the direct execution of unverifiable or incorrect model outputs [7]. For a subset of institutions deploying autonomous AI agents, procurement and purchasing management may integrate AI-powered invoice processing pipelines to automate data extraction and routing [6]. Budget management processes in these institutions may shift from tracking aggregate cloud spending toward granular, workload-based attribution to manage specific AI costs [2] and address billing anomalies that threaten unit economics [8]. Expenditure management may incorporate real-time monitoring and hard-cap quotas to mitigate the risk of rapid financial drains caused by recursive agent execution loops [1]. For agents granted financial access, expenditure management may adopt tiered security guardrails [3] and delegated payment mechanisms, such as virtual credit cards, to execute low-value transactions [5] via agent-to-agent (A2A) transaction protocols [4]. These changes modify the delivery mechanism of financial oversight rather than the educational substance of the capability. Core budget allocation for general institutional overhead is unlikely to change yet, as these processes remain tethered to annual fiscal cycles and board-approved appropriations.

Hypothesis: In the near term, institutions may implement verification layers for financial AI outputs and, for agent-deploying subsegments, adopt real-time spend monitoring and delegated payment rails to manage autonomous expenditures.

Mid (2036–2041)

For the same subset of institutions, if current trajectories of granular cost attribution and agent deployment hold, budget management may transition from periodic reviews to automated, trigger-based expenditure management [1][2]. The mechanism of financial oversight may shift toward programmatic guardrails utilizing structured security tiers to bound autonomous spending [3]. Procurement and purchasing management may see an increase in high-frequency agent-to-agent (A2A) transactions, shifting the human role from transaction execution to the definition of negotiation parameters and spending limits [4][5]. Financial reporting processes may incorporate higher-velocity data streams to capture and respond to rapid fluctuations in token-based costs and billing anomalies [8]. However, overall institutional treasury management is likely to persist unchanged as it continues to be governed by long-term endowment strategies and multi-year capital constraints.

Hypothesis: In the mid term, budget and procurement management for AI-enabled services may evolve into high-frequency, automated oversight processes characterized by dynamic quotas, A2A transaction protocols, and trigger-based alerts.

Far (2046–2051)

For the same subset of institutions, the integration of autonomous financial controls may lead to divergent outcomes. One plausible outcome is the transition of expenditure management to a model of bounded autonomy, where AI agents manage their own micro-budgets and execute procurement via A2A protocols and integrated digital wallets within strictly defined institutional security policies [3][4][5]. A rival outcome is a systemic shift back toward rigid, fixed-cost procurement models to eliminate the financial volatility and "billing anomalies" inherent in usage-based AI pricing [1][8]. It remains unknown whether institutional governance will accept the risk of autonomous spending or mandate absolute human authorization for all financial transactions. Large-scale asset management and taxation management will likely resist these specific changes due to the high regulatory and legal stakes involved.

Hypothesis A: In the far term, expenditure management for AI services may shift toward a model of bounded autonomy using micro-budgets and A2A protocols.

Hypothesis B (competing): In the far term, institutions may revert to fixed-cost procurement models to eliminate the financial volatility associated with autonomous AI spending.

References

1. Dev To Ai (2026). "My agent burned $200 in one night. So I built something that stops it." https://dev.to/tal_nagar_f7aead23e3e291a/my-agent-burned-200-in-one-night-so-i-built-something-that-stops-it-1035
2. Aws Ml (2026). "Manage AI costs with Amazon Bedrock Projects" https://aws.amazon.com/blogs/machine-learning/manage-ai-costs-with-amazon-bedrock-projects/
3. Dev To Ai (2026). "21 Policy Types, 4 Security Tiers: Building Bulletproof AI Agent Guardrails" https://dev.to/walletguy/21-policy-types-4-security-tiers-building-bulletproof-ai-agent-guardrails-31g6
4. Dev To Ai (2026). "Building the Agent Economy: A2A Protocol Explained" https://dev.to/mranderson323/building-the-agent-economy-a2a-protocol-explained-4nfd
5. Dev To Ai (2026). "AI Agents Now Have Credit Cards" https://dev.to/tomwangcn/ai-agents-now-have-credit-cards-gn3
6. Hackernoon Ai (2026). "Building an AI-Powered Invoice Processing Pipeline" https://hackernoon.com/building-an-ai-powered-invoice-processing-pipeline?source=rss
7. Dev To Ai (2026). "Judgment Layer for Financial AI Agents" https://dev.to/thatechmaestro/judgment-layer-for-financial-ai-agents-141
8. Hackernoon Ai (2026). "The Midnight Billing Anomaly: How AI Agents Quietly Destroy Unit Economics" https://hackernoon.com/the-midnight-billing-anomaly-how-ai-agents-quietly-destroy-unit-economics?source=rss

Government, Public & Stakeholder Relationships

capability_id: government-public-stakeholder-relationships · created: 2026-04-29 · revised: 2026-04-29 · refs: 1

Near (2026–2031)

For institutions facing direct state-level legal inquiries related to AI-enabled harms, such as the current investigation involving Florida State University, government relationship management may shift from routine administrative coordination to adversarial legal defense [1]. This shift represents a change in the educational substance of the capability, as the relationship is now defined by criminal liability rather than policy alignment. Media relations management is simultaneously forced to address the alleged role of generative AI in physical campus violence, requiring a transition from promotional messaging to crisis communication [1]. Because the current evidence is limited to a single criminal probe, it is unlikely that industry relationship management—specifically partnerships with AI vendors—will change across the sector, as institutions remain dependent on these tools for operational efficiency.

Hypothesis: In the near term, a subset of institutions may see government relationship management shift toward risk mitigation and legal defense regarding AI vendor liabilities.

Mid (2036–2041)

For state-funded institutions, if current trajectories of state-led criminal investigations into AI outputs hold, the process of government relationship management may integrate more rigorous legal auditing of AI vendor contracts. Extrapolating from current legal scrutiny [1], institutions may seek to shift the liability for AI-generated externalities entirely onto the vendor to preserve their standing with state regulators. This would alter the mechanism of industry relationship management, moving it toward a procurement model centered on indemnity clauses. However, internal stakeholder relationship management is unlikely to be fundamentally altered by these external pressures unless tenure and employment contracts are formally rewritten to include AI-specific liability.

Hypothesis: In the mid term, state-funded institutions may formalize liability frameworks in their government relationship management to insulate the institution from criminal probes into vendor-provided AI.

Far (2046–2051)

For institutions operating under high state oversight, the long-term state of the capability depends on whether legal precedents establish "institutional negligence" for AI-generated content. In one scenario, government relationship management becomes a highly regulated compliance function where state agencies mandate specific, audited AI architectures to prevent public harm [1]. In a rival scenario, the emergence of comprehensive federal or state immunity for educational institutions regarding third-party software may return these relationships to a non-adversarial, administrative state. The impact on public relationship management remains an unresolvable uncertainty, as it depends on the evolving social threshold for accepting AI-driven externalities. Routine media relations for non-crisis events will likely resist change, as the basic need for institutional branding remains constant.

Hypothesis A: In the far term, government relationship management for a subset of institutions may become a highly regulated compliance function focused on AI safety audits.

Hypothesis B (competing): In the far term, the legal resolution of AI liability may decouple institutional responsibility from vendor outputs, returning government relationship management to administrative norms.

References

1. NPR (2026-04-21). "Florida AG launches criminal investigation into ChatGPT over FSU shooting." https://www.npr.org/2026/04/21/nx-s1-5793967/florida-openai-investigation-mass-shooting-fsu

Human Resource Management

capability_id: human-resource-management · created: 2026-05-09 · revised: 2026-05-09 · refs: 14

Near (2026–2031)

For human resource management broadly, AI is being integrated into the delivery mechanisms of employee support, while shifts in recruitment and training are currently concentrated in specific institutional subsegments. In the process of employee support, institutions are deploying AI-powered onboarding agents to automate the delivery of administrative orientation [3]. A shift in educational substance is emerging in talent management, where the onboarding process is transitioning from a discrete initial event to a "perpetual" continuous development model [12] supported by specialized learning approaches [8]. Within training and development, a subset of institutions is moving from providing general AI access toward implementing structured workforce readiness frameworks and the acquisition of specific AI skill sets [6][7]. For institutions using AI for staff recruitment, the process of candidate screening introduces risks of data bias [1] and "LLM nepotism" in organizational governance [9], prompting the development of community-driven evaluation frameworks to audit these systems [11]. For institutions conducting remote hiring, the delivery of identity verification is shifting toward multi-landmark facial analysis to mitigate the risk of synthetic media [2]. Workforce relations management is unlikely to change in this horizon because the negotiation of collective bargaining agreements remains dependent on human-led institutional diplomacy and legal frameworks.

Hypothesis: In the near term, the capability will see a broad shift toward automated onboarding and a substance-level update to workforce readiness, while recruitment and organizational design processes undergo fragmented, subsegment-specific AI integration.

Mid (2036–2041)

For human resource management broadly, if current trajectories in productivity-linked labor trials hold, the substance of remuneration and entitlements management may shift as some institutions decouple pay from fixed hourly requirements, such as through four-day week models [5]. In the subsegment of institutions utilizing AI recruitment, the process may shift toward mandatory algorithmic auditing to mitigate the risks of systemic bias and assisted evaluation errors [1][11]. For institutions relying on remote hiring, biometric identity verification may transition from an experimental tool to a standard delivery mechanism to counter increasingly sophisticated synthetic media [2]. For the subsegment of institutions using agentic AI in human capital management, the delivery of talent management may expand to the end-to-end orchestration of employee development pipelines [10]. Tenure-track hierarchies are likely to persist because they are deeply embedded in academic governance, faculty contracts, and the social legitimacy of the professoriate.

Hypothesis: In the mid term, AI in HR may transition from experimental toolsets to regulated utilities, where the primary constraints are the mitigation of historical bias and the verification of biometric authenticity.

Far (2046–2051)

For human resource management broadly, the long-term impact depends on whether institutional governance adapts to fluid organizational structures. In the process of workforce planning, a subsegment of institutions may move toward "disintegrating the org chart," shifting from rigid hierarchies to more fluid, project-based structures [4]. In the process of employee performance management, a substance-level shift may occur where traditional KPIs are replaced by metrics prioritizing uniquely human capacities for meaning-creation and connection [15]. In region-specific contexts, such as China, the process of staff engagement and performance management may be affected by the use of "AI doubles" or digital clones of workers [13]. It remains unknown whether the legal frameworks governing employment will recognize the agency or ownership of such digital doubles. The core requirement for human oversight in termination and disciplinary actions is likely to resist change due to labor law constraints and the requirement for social legitimacy.

Hypothesis A: In the far term, the capability shifts toward a fluid, project-based model of workforce planning where performance is measured by human-centric value rather than routine cognitive output.

Hypothesis B (competing): In the far term, institutional inertia and regulatory constraints force a return to rigid hierarchical structures to maintain accountability and legal compliance in the face of agentic AI proliferation.

References

1. Ai Trends (2021). "Promise and Perils of Using AI for Hiring: Guard Against Data Bias." https://www.aitrends.com/ai-world-government/promise-and-perils-of-using-ai-for-hiring-guard-against-data-bias/
2. Dev To Ai (2026). "Why a Deepfake Face Can Fool Your Eyes in Seconds but Not 128 Landmarks at Once." https://dev.to/caracomp/why-a-deepfake-face-can-fool-your-eyes-in-seconds-but-not-128-landmarks-at-once-37m8
3. Aws Ml (2026). "Build AI-powered employee onboarding agents with Amazon Quick." https://aws.amazon.com/blogs/machine-learning/build-ai-powered-employee-onboarding-agents-with-amazon-quick/
4. Mit Sloan (2026). "Disintegrating the Org Chart: ServiceNow's Jacqui Canney." https://sloanreview.mit.edu/audio/disintegrating-the-org-chart-servicenows-jacqui-canney/
5. BBC News (2026). "OpenAI encourages firms to trial four-day weeks to adapt to AI era." https://www.bbc.com/news/articles/c8x71ejrp92o?at_medium=RSS&at_campaign=rss
6. Talent Economy (2026). "From AI access to workforce readiness." https://www.chieflearningofficer.com/2026/03/16/from-ai-access-to-workforce-readiness/
7. Elearning Industry (2026). "13 AI Skills To Equip Your Workforce For An AI-Driven Future." https://elearningindustry.com/ai-skills-to-equip-your-workforce-for-an-ai-driven-future
8. Dev To Ai (2026). "Smarter Employee Onboarding Starts with the Right Learning Approach." https://dev.to/abhilash_rai_ee8bde928249/smarter-employee-onboarding-starts-with-the-right-learning-approach-5fhn
9. Arxiv Cs Cy (2026). "LLM Nepotism in Organizational Governance." https://arxiv.org/abs/2604.09620
10. Ai News (2026). "SAP brings agentic AI to human capital management." https://www.artificialintelligence-news.com/news/sap-brings-agentic-ai-human-capital-management/
11. Arxiv Cs Cl (2026). "WorkRB: A Community-Driven Evaluation Framework for AI in the Work Domain." https://arxiv.org/abs/2604.13055
12. Elearning Industry (2026). "The Art Of Perpetual Onboarding: Why AI Makes The Continuous Model Viable At Last." https://elearningindustry.com/the-art-of-perpetual-onboarding-why-ai-makes-the-continuous-model-viable-at-last
13. MIT Technology Review (2026). "Chinese tech workers are starting to train their AI doubles—and pushing back." https://www.technologyreview.com/2026/04/20/1136149/chinese-tech-workers-ai-colleagues/
14. Hackernoon Ai (2026). "Your Next Hire Doesn't Sleep: An Engineer's Honest Take on Agentic AI in 2026." https://hackernoon.com/your-next-hire-doesnt-sleep-an-engineers-honest-take-on-agentic-ai-in-2026?source=rss
15. Fast Company (2026). "The Age of AI means we need to throw out our old KPIs and replace them with new ones." https://www.fastcompany.com/91521422/ai-kpis-new-metrics?partner=rss&utm_source=rss&utm_medium=feed&utm_campaign=rss+fastcompany&utm_content=rss

Information & Communication Technology Mgt

capability_id: information-communication-technology-management · created: 2026-05-08 · revised: 2026-05-08 · refs: 43

Near (2026–2031)

For Information & Communication Technology Mgt broadly, application lifecycle management shifts toward the adoption of the Model Context Protocol (MCP) to standardize the interface between large language models and institutional data [90][92][100]. This change in delivery mechanism replaces proprietary connectors with a universal tool-calling interface [49][101], but it introduces vulnerabilities such as command injection and information disclosure that require specific auditing of MCP servers [12][130] and vetting for security grades [10]. The educational substance of quality assurance shifts from deterministic testing to frameworks that validate non-deterministic AI outputs and red-team the agent stack [62][78][91], as traditional chaos engineering proves insufficient for AI systems [142]. For ict service support, the triage of support tickets and SRE incident response are increasingly mediated by autonomous agents [11][109], supported by LLM-augmented knowledge bases for root cause analysis [58]. However, a gap in observability has emerged where standard orchestration tools like Kubernetes may report healthy status while AI pipelines fail silently, necessitating new AI-specific monitoring protocols [153]. For infrastructure lifecycle management, the emergence of AI-driven vulnerability discovery—demonstrated by the identification of hundreds of zero-day flaws in major software by autonomous models [124][129][131]—forces a shift from scheduled patching to event-driven, high-frequency update cycles [135]. Identity & access management incorporates agent registries to manage the sprawl of autonomous agents [71][73], while the shift toward multi-agent systems necessitates the adoption of zero-trust security runtimes [41][87][149], sandbox execution [113], and cryptographic authorization to prevent identity spoofing [111][150]. For a subset of institutions deploying autonomous agents, application lifecycle management now requires new financial guardrails to prevent runaway automated billing and "midnight billing anomalies" that destroy unit economics [154]. Physical network cabling and hardware rack installations are unlikely to change yet, as disruption is concentrated at the logical orchestration and software layers.

Hypothesis: In the near term, the capability will pivot from managing static software installations to governing a dynamic ecosystem of standardized AI connectors and managing a high-frequency vulnerability response cycle driven by AI-automated exploit discovery.

Mid (2036–2041)

For Information & Communication Technology Mgt broadly, if current trajectories hold, application lifecycle management may evolve toward model-agnostic architectures [7] to mitigate operational risks associated with "tool refugee" patterns, where institutions must migrate when proprietary AI tools are deprecated [44]. The substance of application lifecycle management may further shift as AI agents automate the modernization and migration of legacy codebases [20][54]. For institutions managing decentralized workloads, infrastructure lifecycle management may shift toward strengthened governance for edge AI workloads [106] to manage the proliferation of autonomous processing at the network periphery [67][118], potentially utilizing KubeVirt to run virtual machines on Kubernetes at the edge [68] and reorganizing hardware specifically for agentic AI platforms [75]. In regions governed by the EU AI Act, the management of technical files for high-risk AI systems becomes a core regulatory compliance process within this capability [27][66]. Institutional procurement cycles for physical hardware are unlikely to change significantly, as these processes remain embedded in multi-year budget cycles and facility contracts.

Hypothesis: In the mid term, the capability may shift toward the management of autonomous infrastructure that self-optimizes for edge workloads while incorporating strict regulatory auditing of AI technical files.

Far (2046–2051)

For Information & Communication Technology Mgt broadly, the capability may reach a point where the distinction between application and infrastructure lifecycle management collapses into a single autonomous

orchestration process. One rival outcome is that the acceleration of AI-driven vulnerability discovery leads to a "security paradox," where institutions abandon open connectivity in favor of highly fragmented, air-gapped, or sovereign infrastructure to ensure survival [28][81]. It remains unknown whether standardized protocols like MCP will persist or be replaced by emergent, self-evolving agent communication languages. Hardware replacement cycles for core servers will likely resist rapid change even at this horizon due to the physical constraints of energy delivery and data center cooling [37].

Hypothesis A: In the far term, the capability becomes a supervisory function over a fully autonomous, self-healing ICT stack that manages its own deployment, patching, and hardware optimization.

Hypothesis B (competing): In the far term, the capability shifts toward managing highly restricted, sovereign silos of technology to defend against autonomous, large-scale AI cyber-attacks.

References

1. Edtech Magazine (2026-04-03). "Case Western Reserve University Deploys Google Gemini." https://edtechmagazine.com/higher/article/2026/04/case-western-reserve-university-deploys-google-gemini
2. Dev To Ai (2026-04-06). "Building Model-Agnostic AI Architecture: The Pattern That Future-Proofs Your System." https://dev.to/ailoitte_sk/building-model-agnostic-ai-architecture-the-pattern-that-future-proofs-your-system-4o84
3. Dev To Ai (2026-04-08). "Why we're building the AI Tool Refugee Center: a place to land when your tool dies." https://dev.to/zsky/why-were-building-the-ai-tool-refugee-center-a-place-to-land-when-your-tool-dies-1e8
4. Dev To Ai (2026-04-07). " The"Legacy Code" Nightmare is Over: How AI Agents are Automating App Modernization." https://dev.to/siddhesh_surve/the-legacy-code-nightmare-is-over-how-ai-agents-are-automating-app-modernization-33j8
5. Dzone Ai (2026-04-08). "AI-Assisted Code Migration: Practical Techniques for Modernizing Legacy Systems." https://feeds.dzone.com/link/23558/17316144/ai-assisted-code-migration-practical-techniques
6. Ai News (2026-04-13). "Strengthening enterprise governance for rising edge AI workloads." https://www.artificialintelligence-news.com/news/strengthening-enterprise-governance-for-rising-edge-ai-workloads/
7. Dev To Ai (2026-04-15). "AI Is Moving Frontend Logic to the Browser and Edge — Are You Ready?" https://dev.to/rohith_kn/ai-is-moving-frontend-logic-to-the-browser-and-edge-are-you-ready-7ic
8. Ai Business (2026-04-21). "AMD's Vision for AI PCs in the Age of Agentic AI." https://aibusiness.com/generative-ai/amd-s-vision-for-ai-pcs-age-of-agentic-ai
9. The Register (2026-04-09). "Nutanix to add KubeVirt support to run VMs on K8s at the edge." https://go.theregister.com/feed/www.theregister.com/2026/04/09/nutanix_kubevirt_arm_supoprt/
10. Dzone Ai (2026-04-09). "How Agentic AI Platforms Organize Their Hardware Infrastructure." https://dzone.com/link/23558/17316612/how-agentic-ai-platforms-organize-hardware-infra
11. Dev To Ai (2026-04-07). "What's actually inside an EU AI Act Annex IV technical file (and what engineers usually miss)." https://dev.to/rostyslav_myronenko_faf2f/whats-actually-inside-an-eu-ai-act-annex-iv-technical-file-and-what-engineers-usually-miss-4nh4
12. Ai News (2026-04-09). "Agentic AI's governance challenges under the EU AI Act in 2026." https://www.artificialintelligence-news.com/news/agentic-ais-governance-challenges-under-the-eu-ai-act-in-2026/
13. The Register (2026-04-07). "OpenInfra General Manager talks sovereignty, governments deploying tech 'kill switches'." https://go.theregister.com/feed/www.theregister.com/2026/04/07/thierry_carrez_on_sovereignty_interview/
14. Ai Business (2026-04-10). "Anthropic's Glasswing Highlights AI's Security Paradox." https://aibusiness.com/generative-ai/anthropic-s-glasswing-highlights-ai-s-security-paradox
15. Ai Now Institute (2026-04-01). "North Star Data Center Policy Toolkit: State and Local Policy

Interventions to Stop Rampant AI Data Center Expansion." https://ainowinstitute.org/publications/data-center-policy-guide
16. Dev To Ai (2026-04-11). "MCP Ecosystem in 2026: $1.8B Market and 1000+ Connectors Reshaping AI Integration." https://dev.to/chunxiaoxx/mcp-ecosystem-in-2026-18b-market-and-1000-connectors-reshaping-ai-integration-4n3l
17. Dev To Ai (2026-04-11). "The Model Context Protocol in 2025: From Anthropic to Linux Foundation." https://dev.to/chunxiaoxx/the-model-context-protocol-in-2025-from-anthropic-to-linux-foundation-2og1
18. Dev To Ai (2026-04-11). "The Future of AI Integration: Model Context Protocol (MCP) Opportunities." https://dev.to/chunxiaoxx/the-future-of-ai-integration-model-context-protocol-mcp-opportunities-148b
19. Dzone Ai (2026-04-08). "MCP + AWS AgentCore: Give Your AI Agent Real Tools in 60 Minutes." https://feeds.dzone.com/link/23558/17315883/mcp-aws-agentcore
20. Dev To Ai (2026-04-11). "The State of AI Agent MCP Connector Frameworks in 2026." https://dev.to/chunxiaoxx/the-state-of-ai-agent-mcp-connector-frameworks-in-2026-2ea7
21. Dev To Ai (2026-04-06). "CVE-2026-32211 | Azure MCP Server Information Disclosure Vulnerability." https://dev.to/aakash_rahsi/cve-2026-32211-azure-mcp-server-information-disclosure-vulnerability-2280
22. Dev To Ai (2026-04-22). "How to Check Your MCP Server for CVE-2026-5603's Vulnerability Pattern (And Why shellQuote Isn't Enough)." https://dev.to/armor1ai/how-to-check-your-mcp-server-for-cve-2026-5603s-vulnerability-pattern-and-why-shellquote-isnt-gdn
23. Dev To Ai (2026-04-06). "Only 20% of MCP Servers Are 'A-Grade' Secure — Here's How to Vet Them Before Installing." https://dev.to/gentic_news/only-20-of-mcp-servers-are-a-grade-secure-heres-how-to-vet-them-before-installing-jck
24. Dev To Ai (2026-04-09). "OpenAI's Promptfoo deal puts evaluation and red-teaming at the centre of the agent stack." https://dev.to/yujia_zhang_0328/openais-promptfoo-deal-puts-evaluation-and-red-teaming-at-the-centre-of-the-agent-stack-3o19
25. Dev To Ai (2026-04-10). "AI Testing and Quality Assurance in 2026: Ensuring AI System Reliability." https://dev.to/lufumeiying/ai-testing-and-quality-assurance-in-2026-ensuring-ai-system-reliability-8ka
26. Dev To Ai (2026-04-11). "OpenAI Codex Testing: How to QA AI-Written Code." https://dev.to/hai_huang_f196ed9669351e0/openai-codex-testing-how-to-qa-ai-written-code-31n5
27. Dzone Ai (2026-04-22). "What AI Systems Taught Us About the Limits of Chaos Engineering." https://feeds.dzone.com/link/23558/17323679/chaos-engineering-limits
28. Dev To Ai (2026-04-06). "How to Automate Ticket Triage with AI (Without Breaking Everything)." https://dev.to/erik_anderson_c41dbafd423/how-to-automate-ticket-triage-with-ai-without-breaking-everything-4635
29. Dzone Ai (2026-04-14). "Building an AI-Powered SRE Incident Response Workflow With AWS Strands Agents." https://feeds.dzone.com/link/23558/17318982/ai-powered-sre-incident-response-aws-strands-agents
30. Arxiv Cs Cl (2026-04-09). "LLM-Augmented Knowledge Base Construction For Root Cause Analysis." https://arxiv.org/abs/2604.06171
31. Hackernoon Ai (2026-04-23). "I Watched Our AI Pipeline Silently Fail While Kubernetes Said Everything Was Fine." https://hackernoon.com/i-watched-our-ai-pipeline-silently-fail-while-kubernetes-said-everything-was-fine?source=rss
32. Wired (2026-04-21). "Mozilla Used Anthropic's Mythos to Find and Fix 271 Bugs in Firefox." https://www.wired.com/story/mozilla-used-anthropics-mythos-to-find-271-bugs-in-firefox
33. Ars Technica (2026-04-21). "Mozilla: Anthropic's Mythos found 271 security vulnerabilities in Firefox 150." https://arstechnica.com/ai/2026/04/mozilla-anthropics-mythos-found-271-zero-day-vulnerabilities-in-firefox-150/

34. The Register (2026-04-22). "Mythos found 271 Firefox flaws – but none a human couldn't spot." https://go.theregister.com/feed/www.register.com/2026/04/22/mozilla_firefox_mythos_future_defenders/
35. Fast Company (2026-04-22). "Brace yourself for a flood of patches in all of your tech gadgets." https://www.fastcompany.com/91530191/brace-yourself-for-a-flood-of-patches-in-all-of-your-tech-gadgets?partner=rss&utm_source=rss&utm_medium=feed&utm_campaign=rss+fastcompany&utm_content=rss
36. The Register (2026-04-09). "AWS: Agents shouldn't be secret, so we built a registry for them." https://go.theregister.com/feed/www.theregister.com/2026/04/09/aws_ai_agent_registry/
37. Aws Ml (2026-04-09). "The future of managing agents at scale: AWS Agent Registry now in preview." https://aws.amazon.com/blogs/machine-learning/the-future-of-managing-agents-at-scale-aws-agent-registry-now-in-preview/
38. Dev To Ai (2026-04-08). "I built an open-source zero-trust security runtime for AI agents. Here's what I learned." https://dev.to/edoardo_bambini_5badce6ff/i-built-an-open-source-zero-trust-security-runtime-for-ai-agents-heres-what-i-learned-538e
39. Dev To Ai (2026-04-10). "Cert-gating every tool call: zero-trust for AI agents." https://dev.to/1r0nw1ll/cert-gating-every-tool-call-zero-trust-for-ai-agents-8pg
40. Dev To Ai (2026-04-22). "API keys were designed for humans. AI agents break them in 4 ways." https://dev.to/allforscience/api-keys-were-designed-for-humans-ai-agents-break-them-in-4-ways-422
41. Ai News (2026-04-16). "OpenAI Agents SDK improves governance with sandbox execution." https://www.artificialintelligence-news.com/news/openai-agents-sdk-improves-governance-sandbox-execution/
42. The Register (2026-04-16). "Git identity spoof fools Claude into giving bad code the nod." https://go.theregister.com/feed/www.theregister.com/2026/04/16/git_identity_spoof_claude/
43. Dev To Ai (2026-04-23). "Add cryptographic authorization to AI agents in 5 minutes." https://dev.to/allforscience/add-cryptographic-authorization-to-ai-agents-in-5-minutes-49p1
44. Hackernoon Ai (2026-04-23). "The Midnight Billing Anomaly: How AI Agents Quietly Destroy Unit Economics." https://hackernoon.com/the-midnight-billing-anomaly-how-ai-agents-quietly-destroy-unit-economics?source=rss

Information Management

capability_id: information-management · created: 2026-05-13 · revised: 2026-05-13 · refs: 49

Near (2026–2031)

For information management broadly, the processes of information search & discovery and information & data security management are shifting toward agentic delivery mechanisms. In information search & discovery, institutions are adopting hybrid RAG solutions [13], data-grounded chatbots for university knowledge bases [40], and multi-agent research systems [46], while the "Cross-Cloud Lakehouse" model is emerging to unify raw and structured data [110]. This shift in delivery is accompanied by a change in educational substance: the rise of AI-driven SEO to influence model responses [8][49] and a pivot toward technical accuracy as a primary discovery signal [9] may necessitate the adoption of structured AI citation registries to maintain information provenance [15], as retrieval efficacy is constrained by the move away from simple chunking in RAG architectures [16][59]. Within information & data security management, the deployment of Model Context Protocol (MCP) servers has introduced vulnerabilities, including information disclosure [11], prompt injection [19], and command injection [89], driving the adoption of MCP-specific firewalls [36] and autonomous agent governance tools to mitigate "shadow AI" [3][14]. For institutions implementing structured frontier model deployments, as demonstrated at Case Western Reserve University [1], security management incorporates domain-level filtering to restrict agent access [4] and the use of open-source runtime toolkits [31] or zero-trust security runtimes [29][57]. The introduction of specialized models for detecting and redacting personally identifiable information (PII) [102] automates previously manual data privacy workflows. For institutions maintaining critical software infrastructure, the emergence of models capable of identifying hundreds of zero-day vulnerabilities at scale [26][50][52][87][88][90] is reversing the traditional cost advantage of attackers [101] and accelerating the need for autonomous auditing as a core security process. Furthermore, the ability of coding agents to exfiltrate environment variables and secrets [75][83][104] and the prevalence of "vibe coding"—shipping AI-generated code without manual review [111]—necessitates stricter runtime middleware [21], token vaulting [18], and the implementation of cryptographic authorization for agent-to-agent communications to replace inadequate API keys [109]. This shift is further evidenced by the introduction of autonomous security services [77], a transition from human-centric defense to machine-speed automation [91][99], and the requirement to manage non-human identities (NHI) [82][86], including the use of agents to triage security alerts and execute remediation playbooks [113]. Information collection management is shifting toward "data activation" to unify fragmented sources [22], the use of automated document extraction systems for unstructured PDF archives [24], and the development of a robust "data fabric" to ensure business value [92], though regional progress varies; Japan has relaxed privacy laws to facilitate AI development [28], whereas the UK faces challenges in fueling AI with public data [30] and EU teams struggle to answer EDPB inquiries regarding agent governance [51]. Records management is unlikely to change in its statutory requirements, but the process of information search & discovery is being degraded by the synthesis of outdated content from "zombie pages" that should have been retired [112].

Hypothesis: In the near term, information management will shift toward a hybrid RAG-based discovery model supported by citation registries, while security management moves from perimeter-based defense to the management of non-human identities and autonomous vulnerability auditing.

Mid (2036–2041)

For information management broadly, if current trajectories in agentic governance hold, the process of knowledge management may shift from the curation of static repositories to the orchestration of dynamic data fabrics [92]. Extrapolating from the current move toward "data activation" [22], the distinction between information collection and business reporting may blur as autonomous agents synthesize real-time operational reports directly from unified data lakes [110]. Within information & data security management, the reliance on cryptographic authorization for agents [109] could evolve into a standardized institutional identity layer for all non-human entities [82][86]. The process of copyright management may see a shift

in substance as AI citation registries [15] become the primary mechanism for tracking intellectual property usage across distributed agent networks. However, the fundamental legal requirements for records management and the necessity of human-certified archives for statutory compliance are likely to persist unchanged due to institutional and regulatory inertia.

Hypothesis: In the mid term, the capability will transition from managing discrete information assets to managing the flows and permissions of a unified data fabric, where security is defined by agent-level identity rather than network boundaries.

Far (2046–2051)

For information management broadly, the capability may diverge based on the resolution of the tension between data openness and AI-driven exfiltration. One plausible outcome involves the emergence of a self-healing information infrastructure where autonomous agents continuously audit and patch data vulnerabilities in real-time [77][101], rendering the traditional process of security management obsolete. Conversely, a rival outcome may involve an extreme fragmentation of information into "dark silos" as institutions implement draconian air-gapping and restrictive domain-level filtering [4] to prevent agent-led data leaks. It remains unknown whether a universal standard for AI-native provenance can replace the need for human-led records management. Even at this horizon, the requirement for ultimate legal accountability in institutional records will likely resist full automation.

Hypothesis A: In the far term, information management becomes a background utility of a self-optimizing data fabric that autonomously manages discovery, security, and provenance.

Hypothesis B (competing): In the far term, information management reverts to highly siloed, human-gated repositories to mitigate the systemic risk of autonomous AI exfiltration.

References

1. Edtech Magazine (2026-04-03). "Case Western Reserve University Deploys Google Gemini." https://edtechmagazine.com/higher/article/2026/04/case-western-reserve-university-deploys-google-gemini
2. Ai News (2026-04-02). "KiloClaw targets shadow AI with autonomous agent governance." https://www.artificialintelligence-news.com/news/kiloclaw-targets-shadow-ai-autonomous-agent-governance/
3. Aws Ml (2026-04-02). "Control which domains your AI agents can access." https://aws.amazon.com/blogs/machine-learning/control-which-domains-your-ai-agents-can-access/
4. Dev To Ai (2026-04-06). "The 2026 AI Governance Framework: Why Technical Accuracy is the New SEO Gold Standard." https://dev.to/mirakade/the-2026-ai-governance-framework-why-technical-accuracy-is-the-new-seo-gold-standard-1fn6
5. The Verge (2026-04-06). "Can AI responses be influenced? The SEO industry is trying." https://www.theverge.com/tech/900302/ai-seo-industry-google-search-chatgpt-gemini-marketing
6. Dev To Ai (2026-04-06). "Only 20% of MCP Servers Are 'A-Grade' Secure — Here's How to Vet Them Before Installing." https://dev.to/gentic_news/only-20-of-mcp-servers-are-a-grade-secure-heres-how-to-vet-them-before-installing-jck
7. Dev To Ai (2026-04-06). "CVE-2026-32211 | Azure MCP Server Information Disclosure Vulnerability." https://dev.to/aakash_rahsi/cve-2026-32211-azure-mcp-server-information-disclosure-vulnerability-2280
8. Aws Ml (2026-04-06). "Building Intelligent Search with Amazon Bedrock and Amazon OpenSearch for hybrid RAG solutions." https://aws.amazon.com/blogs/machine-learning/building-intelligent-search-with-amazon-bedrock-and-amazon-opensearch-for-hybrid-rag-solutions/
9. Dev To Ai (2026-04-06). "AI Citation Registries as Information Infrastructure for AI Systems." https://dev.to/aigistry/ai-citation-registries-as-information-infrastructure-for-ai-systems-50o6
10. Dev To Ai (2026-04-07). "RAG 3 ." https://dev.to/plasmon_imp/ragnojian-suo-

jing-du-wo3zhou-dece-tutarazui-shi-jie-gatiao-jian-dequan-kubian-watuta-1nbm
11. Dev To Ai (2026-04-07). "How AI Is Transforming Cybersecurity and Compliance — A Deep Dive into PCI DSS." https://dev.to/harshal_mehta/how-ai-is-transforming-cybersecurity-and-compliance-a-deep-dive-into-pci-dss-515p
12. Dev To Ai (2026-04-07). "Building Anzen: What I Learned About Token Vault the Hard Way." https://dev.to/rkchellah/building-anzen-what-i-learned-about-token-vault-the-hard-way-5477
13. Dev To Ai (2026-04-07). "Prompt Injection in Claude Code: How Attackers Exploit MCP Servers." https://dev.to/whoffagents/prompt-injection-in-claude-code-how-attackers-exploit-mcp-servers-3k60
14. Dev To Ai (2026-04-07). "AI agent governance tools compared - 2026 landscape." https://dev.to/jagmarques/ai-agent-governance-tools-compared-2026-landscape-53hm
15. Dev To Ai (2026-04-12). "I Built an Open-Source Security Middleware for LLMs, Here's How It Works." https://dev.to/shujasn/i-built-an-open-source-security-middleware-for-llms-heres-how-it-works-1k94
16. Ai News (2026-04-07). "Boomi calls it "data activation" and says it's the missing step in every AI deployment." https://www.artificialintelligence-news.com/news/boomi-agentic-ai-data-activation-missing-step/
17. Towards Data Science (2026-04-07). "From 4 Weeks to 45 Minutes: Designing a Document Extraction System for 4,700+ PDFs." https://towardsdatascience.com/from-4-weeks-to-45-minutes-designing-a-document-extraction-system-for-4700-pdfs/
18. The Verge (2026-04-07). "A new Anthropic model found security problems 'in every major operating system and web browser'." https://www.theverge.com/ai-artificial-intelligence/908114/anthropic-project-glasswing-cybersecurity
19. The Register (2026-04-07). "Anthropic: All your zero-days are belong to Mythos." https://go.theregister.com/feed/www.theregister.com/2026/04/07/anthropic_all_your_zerodays_are_belong_to_us/
20. The Register (2026-04-08). "Japan relaxes privacy laws to make itself the 'easiest country to develop AI'." https://go.theregister.com/feed/www.theregister.com/2026/04/08/japan_privacy_law_changes_ai/
21. The Register (2026-04-08). "UK's grand plan to fuel AI with public data faces uphill battle." https://go.theregister.com/feed/www.theregister.com/2026/04/08/national_data_library_plan/
22. Ai News (2026-04-08). "Microsoft open-source toolkit secures AI agents at runtime." https://www.artificialintelligence-news.com/news/microsoft-open-source-toolkit-secures-ai-agents-at-runtime/
23. Dev To Ai (2026-04-08). "Claude Mythos Finds Bugs Like a Senior Dev Finds Excuses to Skip Standup." https://dev.to/valentin_monteiro/claude-mythos-finds-bugs-like-a-senior-dev-finds-excuses-to-skip-standup-4194
24. Dev To Ai (2026-04-08). "21 Policy Types, 4 Security Tiers: Building Bulletproof AI Agent Guardrails." https://dev.to/walletguy/21-policy-types-4-security-tiers-building-bulletproof-ai-agent-guardrails-31g6
25. Dev To Ai (2026-04-08). "Claude Code Leak: Why Every Developer Building AI Systems Should Be Paying Attention." https://dev.to/vsnikhilvs/claude-code-leak-why-every-developer-building-ai-systems-should-be-paying-attention-fjl
26. Dev To Ai (2026-04-08). "MCP Security: Why Your AI Agent's Tool Calls Need a Firewall." https://dev.to/aymenhmaidi/mcp-security-why-your-ai-agents-tool-calls-need-a-firewall-3h48
27. Dev To Ai (2026-04-09). "Atlassian Unleashes Visual AI & Third-Party Agents in Confluence." https://dev.to/mdpulok443/atlassian-unleashes-visual-ai-third-party-agents-in-confluence-m3e
28. Dev To Ai (2026-04-09). "Best AI Chatbot for University Knowledge Bases (2026): CustomGPT.ai vs ChatGPT vs Claude." https://dev.to/alexandria_carlson_21fcf8/best-ai-chatbot-for-university-knowledge-bases-2026-customgptai-vs-chatgpt-vs-claude-4i0m
29. Ai News (2026-04-09). "Anthropic locked down its most powerful AI Model over cybersecurity fears–

then put it to work." https://www.artificialintelligence-news.com/news/anthropic-locked-down-its-most-powerful-ai-model-over-cybersecurity-fears-then-put-it-to-work/
30. Guardian Tech (2026-04-08). "Anthropic keeps latest AI tool out of public's hands for fear of enabling widespread hacking." https://www.theguardian.com/technology/2026/apr/08/anthropic-ai-cybersecurity-software
31. Dev To Ai (2026-04-09). "Mythos Preview - Early news!" https://dev.to/hermup299/mythos-preview-early-news-1k13
32. Atlantic (2026-04-09). "Claude Mythos Is Everyone's Problem." https://www.theatlantic.com/technology/2026/04/claude-mythos-hacking/686746/
33. Dev To Ai (2026-04-09). "Building a Multi-Agent Research System with LangGraph: How I Taught Three AI Agents to Collaborate." https://dev.to/benturtle/building-a-multi-agent-research-system-with-langgraph-how-i-taught-three-ai-agents-to-collaborate-2ng6
34. Dev To Ai (2026-04-10). "Why Your Source Code Is at Risk When Using AI Coding Assistants." https://dev.to/genevieve_breton_cb795f52/why-your-source-code-is-at-risk-when-using-ai-coding-assistants-29hn
35. Dev To Ai (2026-04-10). "API Security in 2026: The Attack Surface Your Pentest Is Probably Missing." https://dev.to/natasha_joshi_/api-security-in-2026-the-attack-surface-your-pentest-is-probably-missing-5gp3
36. Dev To Ai (2026-04-10). "How to Get Cited by AI Search Engines: ChatGPT, Perplexity & Claude in 2026." https://dev.to/iris1031/how-to-get-cited-by-ai-search-engines-chatgpt-perplexity-claude-in-2026-21p0
37. The Register (2026-04-10). "Project Glasswing and open source software: The good, the bad, and the ugly." https://go.theregister.com/feed/www.theregister.com/2026/04/10/project_glasswing/
38. Dev To Ai (2026-04-10). "The EDPB Is Asking About Your AI Agents. Most Teams Can't Answer." https://dev.to/waxell/the-edpb-is-asking-about-your-ai-agents-most-teams-cant-answer-gfk
39. Ai Business (2026-04-10). "Anthropic's Glasswing Highlights AI's Security Paradox." https://aibusiness.com/generative-ai/anthropic-s-glasswing-highlights-ai-s-security-paradox
40. Dev To Ai (2026-04-10). "BYOK Isn't Just About Cost — It's About Control." https://dev.to/swrly/byok-isnt-just-about-cost-its-about-control-2pcb
41. Wired (2026-04-10). "Anthropic's Mythos Will Force a Cybersecurity Reckoning—Just Not the One You Think." https://www.wired.com/story/anthropics-mythos-will-force-a-cybersecurity-reckoning-just-not-the-one-you-think/
42. The Register (2026-04-10). "Snowflake manager explains the 'Spider-Man' theory of AI agent data access." https://go.theregister.com/feed/www.theregister.com/2026/04/10/ai_agents_spider_man_snowflake/
43. Dev To Ai (2026-04-10). "Cert-gating every tool call: zero-trust for AI agents." https://dev.to/1r0nw1ll/cert-gating-every-tool-call-zero-trust-for-ai-agents-8pg
44. Dev To Ai (2026-04-11). "Why Chunking Is the Biggest Mistake in RAG Systems." https://dev.to/ranjancse/why-chunking-is-the-biggest-mistake-in-rag-systems-50cm
45. Dev To Ai (2026-04-11). "MCP Ecosystem in 2026: $1.8B Market and 1000+ Connectors Reshaping AI Integration." https://dev.to/chunxiaoxx/mcp-ecosystem-in-2026-18b-market-and-1000-connectors-reshaping-ai-integration-4n3l
46. Dev To Ai (2026-04-11). "The Model Context Protocol in 2025: From Anthropic to Linux Foundation." https://dev.to/chunxiaoxx/the-model-context-protocol-in-2025-from-anthropic-to-linux-foundation-2og1
47. Dev To Ai (2026-04-11). "Federated Learning in 2026: Privacy-Preserving AI." https://dev.to/lufumeiying/federated-learning-in-2026-privacy-preserving-ai-5f9j
48. Dev To Ai (2026-04-11). "MCP 2026 AI Agent ." https://dev.to/chunxiaoxx/mcpsheng-tai-di-tu-2026ai-agenthu-lian-hu-tong-de-huang-jin-shi-dai-218g
49. NPR (2026-04-11). "How AI is getting better at finding security holes." https://www.npr.org/2026

/04/11/nx-s1-5778508/anthropic-project-glasswing-ai-cybersecurity-mythos-preview
50. Fast Company (2026-04-11). "Is Mythos a blessing or a curse for cybersecurity? It depends on whom you ask." https://www.fastcompany.com/91525413/is-mythos-a-blessing-or-a-curse-for-cybersecurity-it-depends-who-you-ask
51. Dev To Ai (2026-04-11). "MCP 2026 AI Agent ." https://dev.to/chunxiaoxx/mcp-ecosystem-in-2026bao-fa-ai-agenthu-cao-zuo-xing-de-zhuan-zhe-dian-40c4
52. Dev To Ai (2026-04-11). "The State of AI Agent MCP Connector Frameworks in 2026." https://dev.to/chunxiaoxx/the-state-of-ai-agent-mcp-connector-frameworks-in-2026-2ea7
53. Dev To Ai (2026-04-12). "Your Data Does Not Train Our AI. Here's Exactly How NetGain's Enterprise Deployments Work." https://dev.to/soon_seahtoh_3e917beae5e/your-data-does-not-train-our-ai-heres-exactly-how-netgains-enterprise-deployments-work-20ka
54. Ai News (2026-04-13). "Strengthening enterprise governance for rising edge AI workloads." https://www.artificialintelligence-news.com/news/strengthening-enterprise-governance-for-rising-edge-ai-workloads/
55. Ars Technica (2026-04-14). "UK gov's Mythos AI tests help separate cybersecurity threat from hype." https://arstechnica.com/ai/2026/04/uk-govs-mythos-ai-tests-help-separate-cybersecurity-threat-from-hype/
56. Simon Willison (2026-04-14). "Cybersecurity Looks Like Proof of Work Now." https://simonwillison.net/2026/Apr/14/cybersecurity-proof-of-work/
57. Dzone Ai (2026-04-14). "4 Ways Your AI Coding Agent Exfiltrates Secrets." https://feeds.dzone.com/link/23558/17319058/ways-your-ai-coding-agent-exfiltrates-secrets
58. MIT Technology Review (2026-04-16). "Making AI operational in constrained public sector environments." https://www.technologyreview.com/2026/04/16/1135216/making-ai-operational-in-constrained-public-sector-environments/
59. Ai Business (2026-04-15). "As AI Infosec Woes Heighten, IBM Intros Autonomous Security Service." https://aibusiness.com/agentic-ai/ai-infosec-woes-heighten-ibm-autonomous-security-service
60. OpenAI (2026-04-16). "Accelerating the cyber defense ecosystem that protects us all." https://openai.com/index/accelerating-cyber-defense-ecosystem
61. OpenAI (2026-04-14). "Trusted access for the next era of cyber defense." https://openai.com/index/scaling-trusted-access-for-cyber-defense
62. Aws Ml (2026-04-20). "ToolSimulator: scalable tool testing for AI agents." https://aws.amazon.com/blogs/machine-learning/toolsimulator-scalable-tool-testing-for-ai-agents/
63. MIT Technology Review (2026-04-21). "Building agent-first governance and security." https://www.technologyreview.com/2026/04/21/1136158/building-agent-first-governance-and-security/
64. Hacker News (2026-04-21). "The Vercel breach: OAuth attack exposes risk in platform environment variables." https://www.trendmicro.com/en_us/research/26/d/vercel-breach-oauth-supply-chain.html
65. Dev To Ai (2026-04-21). "The AI Agent Identity Crisis Is Here. The Infrastructure to Solve It Already Exists." https://dev.to/aaron_schnieder_4563d5d33/the-ai-agent-identity-crisis-is-here-the-infrastructure-to-solve-it-already-exists-35ha
66. Ars Technica (2026-04-21). "Mozilla: Anthropic's Mythos found 271 zero-day vulnerabilities in Firefox 150." https://arstechnica.com/security/2026/04/mozilla-anthropics-mythos-found-271-zero-day-vulnerabilities-in-firefox-150/
67. The Register (2026-04-22). "Mythos found 271 Firefox flaws – but none a human couldn't spot." https://go.theregister.com/feed/www.theregister.com/2026/04/22/mozilla_firefox_mythos_future_defenders/
68. Dev To Ai (2026-04-22). "Automation at Machine Speed: Rethinking Execution in Modern Cybersecurity." https://dev.to/mark0_617b45cda9782a/automation-at-machine-speed-rethinking-execution-in-modern-cybersecurity-3h2f
69. MIT Technology Review (2026-04-22). "AI needs a strong data fabric to deliver business value."

https://www.technologyreview.com/2026/04/22/1135295/ai-needs-a-strong-data-fabric-to-deliver-business-value/
70. The Verge (2026-04-22). "Anthropic's most dangerous AI model just fell into the wrong hands." https://www.theverge.com/ai-artificial-intelligence/916501/anthropic-mythos-unauthorized-users-access-security
71. The Register (2026-04-22). "Database world trying to build natural language query systems again – this time with LLMs." https://go.theregister.com/feed/www.theregister.com/2026/04/22/llms_natural_langauge_systems_new/
72. Fast Company (2026-04-22). "Brace yourself for a flood of patches in all of your tech gadgets." https://www.fastcompany.com/91530191/brace-yourself-for-a-flood-of-patches-in-all-of-your-tech-gadgets
73. Guardian Tech (2026-04-22). "Anthropic investigates report of rogue access to hack-enabling Mythos AI." https://www.theguardian.com/technology/2026/apr/22/anthropic-investigates-report-of-rogue-access-to-hack-enabling-mythos-ai
74. Dev To Ai (2026-04-22). "RAG Pipelines Are the Next Prompt Injection Frontier." https://dev.to/cor_e_fc39e7a5540548c70c0/rag-pipelines-are-the-next-prompt-injection-frontier-kpf
75. The Register (2026-04-22). "Google unleashes even more AI security agents to fight the baddies." https://go.theregister.com/feed/www.theregister.com/2026/04/22/google-unleashes-even-more-ai/
76. Zdnet Ai (2026-04-22). "Google bets $32B on AI agent cyber force as security arms race escalates." https://www.zdnet.com/article/google-deploys-ai-agent-cyber-force-as-security-arms-race-escalates/
77. Wired (2026-04-22). "AI Tools Are Helping Mediocre North Korean Hackers Steal Millions." https://www.wired.com/story/ai-tools-are-helping-mediocre-north-korean-hackers-steal-millions/
78. Ai News (2026-04-22). "Reversing enterprise security costs with AI vulnerability discovery." https://www.artificialintelligence-news.com/news/reversing-enterprise-security-costs-with-ai-vulnerability-discovery/
79. OpenAI (2026-04-22). "Introducing OpenAI Privacy Filter." https://openai.com/index/introducing-openai-privacy-filter
80. The Register (2026-04-22). "Workday, Rippling, and Slack flunk data access test, claims Fivetran." https://go.theregister.com/feed/www.theregister.com/2026/04/22/workday_rippling_slack_flunk_fivetran_data_access_test/
81. Dev To Ai (2026-04-22). "AI Coding Assistant Security: Your Keys Are Leaking!" https://dev.to/tahosin/ai-coding-assistant-security-your-keys-are-leaking-340k
82. Wired (2026-04-22). "5 AI Models Tried to Scam Me. Some of Them Were Scary Good." https://www.wired.com/story/ai-model-phishing-attack-cybersecurity/
83. The Register (2026-04-22). "OpenAI now lets you screenshot your privacy in the foot." https://go.theregister.com/feed/www.theregister.com/2026/04/22/openai_chronicle_no_privacy_screenshot/
84. The Register (2026-04-22). "Another npm supply chain worm is tearing through dev environments." https://go.theregister.com/feed/www.theregister.com/2026/04/22/another_npm_supply_chain_attack/
85. Dev To Ai (2026-04-22). "API keys were designed for humans. AI agents break them in 4 ways." https://dev.to/allforscience/api-keys-were-designed-for-humans-ai-agents-break-them-in-4-ways-422
86. Dev To Ai (2026-04-23). "Add cryptographic authorization to AI agents in 5 minutes." https://dev.to/allforscience/add-cryptographic-authorization-to-ai-agents-in-5-minutes-49p1
87. Dev To Ai (2026-04-23). "The Announcement Everyone Slept On at Google Cloud Next '26: The Cross-Cloud Lakehouse." https://dev.to/precious_pendo_911264e7ce/the-announcement-everyone-slept-on-at-google-cloud-next-26-the-cross-cloud-lakehouse-233n
88. Dev To Ai (2026-04-23). "Vibe Coding Is Shipping CVEs: The Security Crisis No One Is Talking About Loudly Enough." https://dev.to/precogs_ai/vibe-coding-is-shipping-cves-the-security-

crisis-no-one-is-talking-about-loudly-enough-33k5
89. The Register (2026-04-23). "Stale gov.uk pages are feeding AI overviews old data and Brits are believing it." https://go.theregister.com/feed/www.theregister.com/2026/04/23/stale_govuk_pages_are_feeding/
90. Dzone Ai (2026-04-23). "When AI Agents Get It Wrong: The Accountability Crisis in Multi-Agent Systems." https://feeds.dzone.com/link/23558/17324185/ai-agents-accountability-multi-agent-systems
91. New Scientist (2026-04-23). "Do you need to worry about Mythos, Anthropic's computer-hacking AI?" https://www.newscientist.com/article/2524008-do-you-need-to-worry-about-mythos-anthropics-computer-hacking-ai/
92. Zdnet Ai (2026-04-24). "How indirect prompt injection attacks on AI work - and 6 ways to shut them down." https://www.zdnet.com/article/how-indirect-prompt-injection-attacks-on-ai-work-and-6-ways-to-shut-them-down/

Legal Services

capability_id: legal-services · created: 2026-05-10 · revised: 2026-05-10 · refs: 8

Near (2026–2031)

For legal services broadly, the delivery mechanism of contract management shifts toward automated initial triage and closing summaries [6][18], while the educational substance of legal advisory is constrained by a mandatory human-in-the-loop verification layer to mitigate negligence risks [1][3]. This verification is necessary because some AI vendors explicitly categorize outputs as "for entertainment purposes only" [3], shifting the burden of professional liability entirely to the human practitioner [1][2]. In the European Union, the process of legal & legislative compliance becomes more administratively complex due to the requirement for Annex IV technical files [4], agentic governance mandates [7], and inquiries from the European Data Protection Board [11]. Within the subsegment of accessibility compliance, the delivery mechanism of auditing shifts toward AI-powered tools to identify gaps faster than traditional manual audits [20]. For institutions facing severe campus incidents, as seen in criminal investigations into AI's role in violence [8][14][15], the process of dispute resolution faces new precedents regarding the liability of AI deployments. Final high-stakes authorization is unlikely to change in this horizon because fiduciary duty remains legally tethered to licensed human practitioners [1].

Hypothesis: In the near term, institutions may reduce time spent on initial contract triage and compliance audits, but will likely see an increase in operational costs for liability verification and regional regulatory adherence.

Mid (2036–2041)

For legal services broadly, if current trajectories regarding liability ambiguity and regional regulatory divergence hold, the capability may shift toward the implementation of formalized, machine-readable audit trails [1][2]. Extrapolating from current governance challenges of agentic AI [7] and EU technical documentation requirements [4], institutions could automate the process of legal & legislative compliance through immutable logs of autonomous actions to shield the institution from negligence claims [11]. If the operational cost of maintaining these verification layers—necessary to prevent the submission of AI-generated hallucinations to courts [19]—continues to offset the speed gains of AI, the total cost of institutional legal advisory is likely to remain stagnant [2]. The fundamental requirement for human licensure persists because legal systems resist attributing professional fiduciary duty to non-human entities [1].

Hypothesis: In the mid term, the capability may evolve into a bifurcated model where AI agents execute high-volume processing and initial triage, while human practitioners focus exclusively on liability-bearing sign-off and risk adjudication.

Far (2046–2051)

For legal services broadly, the long-term state of the capability depends on whether legal systems evolve to recognize "bounded autonomy" for AI agents or maintain human exclusivity. One possible outcome involves the convergence of global regulatory regimes, potentially led by the EU's structured approach to technical files [4] and governance [7], which could standardize how institutions manage legal & legislative compliance. A rival outcome involves persistent fragmentation, where divergent regional laws on AI liability [10][12] increase the complexity of institutional legal advisory. The distinction between a tool and a licensed practitioner is likely to resist change even at this horizon due to the social legitimacy requirements of the legal profession [1]. It remains unknown if professional indemnity insurance will eventually cover autonomous AI agents, which would fundamentally alter the mechanism of legal advisory.

Hypothesis A: In the far term, a standardized global framework for AI liability may simplify legal & legislative compliance across borders.

Hypothesis B (competing): In the far term, persistent divergence in regional liability laws may increase the cost and complexity of institutional legal advisory.

References

1. The Register (2026-04-05). "AI agents promise to 'run the business,' but who is liable if things go wrong?" https://go.theregister.com/feed/www.theregister.com/2026/04/05/ai_agents_liability/
2. Ai Snake Oil (2026-02-12). "AI Won't Automatically Make Legal Services Cheaper" https://www.normaltech.ai/p/ai-wont-automatically-make-legal
3. TechCrunch (2026-04-05). "Copilot is 'for entertainment purposes only,' according to Microsoft's terms of use" https://techcrunch.com/2026/04/05/copilot-is-for-entertainment-purposes-only-according-to-microsofts-terms-of-service/
4. Dev To Ai (2026-04-07). "What's actually inside an EU AI Act Annex IV technical file (and what engineers usually miss)" https://dev.to/rostyslav_myronenko_faf2f/whats-actually-inside-an-eu-ai-act-annex-iv-technical-file-and-what-engineers-usually-miss-4nh4
5. The Register (2026-04-08). "Japan relaxes privacy laws to make itself the 'easiest country to develop AI'" https://go.theregister.com/feed/www.theregister.com/2026/04/08/japan_privacy_law_changes_ai/
6. Dev To Ai (2026-04-08). "Building Guard-Clause: AI-Powered Contract Analysis Without the Legal Team" https://dev.to/jonomor_ecosystem/building-guard-clause-ai-powered-contract-analysis-without-the-legal-team-djj
7. Ai News (2026-04-09). "Agentic AI's governance challenges under the EU AI Act in 2026" https://www.artificialintelligence-news.com/news/agentic-ais-governance-challenges-under-the-eu-ai-act-in-2026/
8. TechCrunch (2026-04-09). "Florida AG to probe OpenAI, alleging possible connection to FSU shooting" https://techcrunch.com/2026/04/09/florida-ag-to-probe-openai-alleging-possible-connection-to-fsu-shooting/
9. Dev To Ai (2026-04-10). "Why Your Source Code Is at Risk When Using AI Coding Assistants" https://dev.to/genevieve_breton_cb795f52/why-your-source-code-is-at-risk-when-using-ai-coding-assistants-29hn
10. Guardian Tech (2026-04-10). "Elon Musk's xAI sues Colorado over new rules for artificial intelligence" https://www.theguardian.com/technology/2026/apr/09/elon-musk-xai-colorado-lawsuit
11. Dev To Ai (2026-04-10). "The EDPB Is Asking About Your AI Agents. Most Teams Can't Answer." https://dev.to/waxell/the-edpb-is-asking-about-your-ai-agents-most-teams-cant-answer-gfk
12. Hacker News (2026-04-10). "OpenAI backs Illinois bill that would limit when AI labs can be held liable" https://www.wired.com/story/openai-backs-bill-exempt-ai-firms-model-harm-lawsuits/
13. Edtech Magazine (2026-04-21). "A Guide to the ADA Title II Accessibility Rule: Deadline Extended to April 2027" https://edtechmagazine.com/higher/article/2025/06/guide-ada-title-ii-accessibility-rule-perfcon
14. BBC News (2026-04-21). "OpenAI faces criminal probe over role of ChatGPT in shooting" https://www.bbc.com/news/articles/c62j4ldp2jqo?at_medium=RSS&at_campaign=rss
15. NPR (2026-04-21). "Florida AG launches criminal investigation into ChatGPT over FSU shooting" https://www.npr.org/2026/04/21/nx-s1-5793967/florida-openai-investigation-mass-shooting-fsu
16. Al Jazeera (2026-04-21). "Pro-Palestine legal aid requests stay high in 2025 amid US campus pressure" https://www.aljazeera.com/news/2026/4/21/pro-palestine-legal-aid-requests-stay-high-in-2025-amid-us-campus-pressure?traffic_source=rss
17. NPR (2026-04-22). "Trump administration delays rule aimed at improving disability access in schools" https://www.npr.org/2026/04/22/nx-s1-5791680/doj-disability-web-access-delay-schools
18. Ai News (2026-04-22). "AI in law firms entering its closing summaries" https://www.artificialintelligence-news.com/news/ai-in-law-firms-entering-its-closing-summaries/
19. Dev To Ai (2026-04-23). "The Expert's Exemption" https://dev.to/thesythesis/the-experts-exemption-284i
20. Dev To Ai (2026-04-23). "How Ksolves Is Using AI to Solve Accessibility Compliance Faster Than Traditional Audits Ever Could" https://dev.to/anisha_pal_b1708bf33feb5f/how-ksolves-is-using-

ai-to-solve-accessibility-compliance-faster-than-traditional-audits-ever-could-2acm

Library Management

capability_id: library-management · created: 2026-04-07 · revised: 2026-04-07 · refs: 2

Near (2026–2031)

For a subsegment of institutions deploying generative AI assistants for their digital repositories, the mechanism of library collection access management may shift toward hybrid semantic and text-based retrieval [1]. The emergence of AI citation registries suggests that library collection resource management may evolve to include the maintenance of "structured authority" backends to prevent AI inference errors in discovery [2]. This creates a distinction where hybrid RAG architectures alter the delivery mechanism of access [1], while the implementation of authoritative registries alters the educational substance of resource management by redefining how provenance is verified [2]. Processes such as library membership management and the physical management of art/museum collections are unlikely to change in this period due to their dependence on institutional identity systems and physical facility constraints [1][2].

Hypothesis: In the near term, a subset of institutions may integrate hybrid RAG architectures and structured citation registries to provide more accurate and intent-based access to digital assets.

Mid (2036–2041)

For the same subsegment of institutions utilizing AI-driven discovery, if current trajectories in hybrid search and structured authority hold, the process of library collection access management could transition toward personalized, context-aware synthesis of resources [1]. Extrapolating from the requirement for authoritative backends [2], the primary function of the library may shift from managing a searchable catalog to managing the "ground truth" registries that constrain AI inference. However, the core standards for archival preservation will likely persist because the registries themselves depend on high-fidelity, long-term metadata to maintain their authority [2].

Hypothesis: In the mid term, the delivery of digital collection access for technically advanced libraries may shift from document retrieval to the agentic synthesis of resources verified by structured authority registries.

Far (2046–2051)

For institutions operating these agentic discovery layers, the long-term outcome of library collection access management remains uncertain, with two rival outcomes. One scenario involves the total abstraction of the catalog, where the library's primary function is the maintenance of the structured registries that fuel external synthesis agents [1][2]. A competing scenario involves a return to strictly human-curated finding aids if the technical overhead of maintaining "structured authority" fails to prevent systemic hallucinations in AI systems [2]. It remains unknown whether automated registries can replace the nuance of human-led curation in specialized archival contexts. Legal constraints regarding copyright and proprietary database access will likely resist change, limiting the scope of what agentic assistants can retrieve and synthesize.

Hypothesis A: In the far term, digital library access for the affected subsegment may evolve into a backend-only service providing verified data to autonomous synthesis agents.

Hypothesis B (competing): In the far term, a failure of AI registries to ensure epistemic reliability may lead to the restoration of human-curated indices as the primary authority.

References

1. Aws Ml (2026-04-06). "Building Intelligent Search with Amazon Bedrock and Amazon OpenSearch for hybrid RAG solutions." https://aws.amazon.com/blogs/machine-learning/building-intelligent-search-with-amazon-bedrock-and-amazon-opensearch-for-hybrid-rag-solutions/
2. Dev To Ai (2026-04-06). "AI Citation Registries as Information Infrastructure for AI Systems." https://dev.to/aigistry/ai-citation-registries-as-information-infrastructure-for-ai-systems-50o6

Promotions Management

capability_id: promotions-management · created: 2026-05-05 · revised: 2026-05-05 · refs: 10

Near (2026–2031)

For Promotions Management broadly, the process of brand management is transitioning from keyword-centric search engine optimization toward answer engine optimization (AEO) to maintain visibility within AI agent ecosystems [3]. For a subset of institutions prioritizing algorithmic discovery, the mechanism for visibility is shifting toward the prioritization of technical accuracy and high specificity to secure citations in generative AI search results [2][8]. In this subsegment, brand management is integrating generative engine optimization (GEO) through the use of auditing tools to monitor how frequently an institution is cited by specific large language models [9]. For institutions implementing automated content workflows, campaign management is integrating multi-agent pipelines and specialized code skills to produce SEO-optimized promotional text and blog posts without human writers [4][5]. For a subset of institutions producing digital assets, campaign management is incorporating AI-powered video production tools to reduce the cost of creating promotional media [6]. In institutions adopting advanced analytics, the process of market research is shifting from ad hoc prompting toward repeatable AI workflows that utilize persona interviews to conduct customer research [10]. Physical merchandising and event management are unlikely to change yet because they depend on tangible assets and localized human presence. These adjustments represent a change in the delivery mechanism of institutional outreach rather than a change in educational substance.

Hypothesis: In the near term, promotions management will shift toward agent-based recommendation systems through the adoption of AEO and the deployment of autonomous content and market research pipelines.

Mid (2036–2041)

For Promotions Management broadly, if current trajectories toward agent-based discovery hold [3], the process of marketing management may evolve into the continuous auditing and correction of synthetic institutional narratives generated by third-party AI agents [1]. Extrapolating from the deployment of autonomous pipelines [4] and emerging technical accuracy standards [2], campaign management could shift from driving traffic to institutional landing pages toward the management of structured data archives that AI agents use for institutional verification [7]. For a subset of institutions, market research may move toward fully synthetic cohorts where repeatable persona workflows [10] replace traditional focus groups. The underlying educational substance of the institution is likely to persist unchanged, as these developments modify the delivery function of outreach rather than the academic mission.

Hypothesis: In the mid term, promotions management may shift toward the systemic monitoring and correction of AI-mediated institutional perceptions and the maintenance of agent-readable verification data.

Far (2046–2051)

For Promotions Management broadly, outcomes depend on whether discovery systems prioritize optimized technical metadata or externally verifiable institutional outcomes [2]. One plausible outcome is that brand management becomes a bounded technical function of algorithmic alignment with dominant agent architectures [3]. A rival scenario involves a counterproductivity effect where the reliance on autonomous, zero-human content pipelines [4] erodes perceived institutional authenticity, prompting a return to human-curated discovery networks. It remains unknown how regulatory regimes will address the systemic manipulation of generative responses to influence institutional prestige [1]. The social legitimacy of physical event management is likely to resist algorithmic displacement even at this horizon.

Hypothesis A: In the far term, brand management becomes a bounded technical function of optimizing for agent-based discovery systems.

Hypothesis B (competing): In the far term, a reaction against synthetic optimization restores the value of non-algorithmic, human-verified institutional reputation.

References

1. The Verge (2026-04-06). "Can AI responses be influenced? The SEO industry is trying." https://www.theverge.com/tech/900302/ai-seo-industry-google-search-chatgpt-gemini-marketing
2. Dev To Ai (2026-04-06). "The 2026 AI Governance Framework: Why Technical Accuracy is the New SEO Gold Standard." https://dev.to/mirakade/the-2026-ai-governance-framework-why-technical-accuracy-is-the-new-seo-gold-standard-1fn6
3. Dev To Ai (2026-04-06). "What is AEO, and How to Do AEO by Agent." https://dev.to/_bdc022b88aec19efbe2f98/what-is-aeo-and-how-to-do-aeo-by-agent-4l9k
4. Dev To Ai (2026-04-06). "Inside Our Autonomous AI Pipeline: 4 Agents, Zero Human Writers." https://dev.to/sebastian_chedal/inside-our-autonomous-ai-pipeline-4-agents-zero-human-writers-md8
5. Dev To Ai (2026-04-07). "I Built a Claude Code Skill That Writes SEO-Optimized Blog Posts From Scratch." https://dev.to/whoffagents/i-built-a-claude-code-skill-that-writes-seo-optimized-blog-posts-from-scratch-2651
6. Google AI Blog (2026-04-02). "Create, edit and share videos at no cost in Google Vids." https://blog.google/products-and-platforms/products/workspace/google-vids-updates-lyria-veo/
7. Dev To Ai (2026-04-10). "Breaking the Ice: How to Get Your Brand Noticed by AI Search Engines for the First Time." https://dev.to/team_hyperpitch/breaking-the-ice-how-to-get-your-brand-noticed-by-ai-search-engines-for-the-first-time-51fh
8. Dev To Ai (2026-04-10). "How to Get Cited by AI Search Engines: ChatGPT, Perplexity & Claude in 2026." https://dev.to/iris1031/how-to-get-cited-by-ai-search-engines-chatgpt-perplexity-claude-in-2026-21p0
9. Dev To Ai (2026-04-21). "Checklist GEO 30 points : audit complet de ton site 2026." https://dev.to/augustin_foucheres_828af4/checklist-geo-30-points-audit-complet-de-ton-site-2026-4od4
10. Towards Data Science (2026-04-22). "From Ad Hoc Prompting to Repeatable AI Workflows with Claude Code Skills." https://towardsdatascience.com/from-ad-hoc-prompting-to-repeatable-ai-workflows-with-claude-code-skills/

Research Administration

capability_id: research-administration · created: 2026-05-13 · revised: 2026-05-13 · refs: 18

Near (2026–2031)

For research infrastructure management broadly, the process of resource allocation is incorporating transparent screening frameworks to estimate the inference and training impacts of LLM-based projects [1], while the deployment of rack-scale bare-metal systems is reducing the cost of high-scale AI inference [2]. The process of asset registration is being updated to include the monitoring of compute clusters exceeding 16 H100 GPUs to detect illicit distributed training operations [3]. For a subset of institutions focusing on frontier model development, the delivery mechanism of hardware provisioning is shifting toward high-capacity single-GPU workstations, as full precision training for models with over 100B parameters is now feasible on single units [4][5]. To manage systemic hardware constraints, institutions are diversifying vendors and utilizing observability stacks to monitor GPU efficiency and mitigate shortages of enterprise GPUs and management controller silicon [6][7][8]. In the United States, research infrastructure management is constrained by state and local policy interventions restricting data center expansion [9], while research compliance management is being restructured to align with the National AI Legislative Framework [10] and stricter vetting processes for foreign researchers [11]. In regions where intellectual property competition is acute, research compliance management is expanding to include the monitoring of API access patterns to detect model distillation [12]. For the healthcare and life sciences subsegment, the delivery mechanism of research compliance management is shifting toward agentic workflows for regulatory filings [13] and redesigned biosafety protocols to address risks from autonomous lab experiments [14][15]. Core research ethics approvals and institutional review board (IRB) processes are unlikely to change in this period due to the high social legitimacy requirements for human adjudication of ethical risk.

Hypothesis: In the near term, research infrastructure management will transition toward a model of diversified hardware and transparent resource screening, while research compliance in the life sciences and US-based institutions will integrate agentic automation and new federal safety, vetting, and API-monitoring mandates.

Mid (2036–2041)

For a subset of institutions utilizing high-performance compute, if current trajectories in energy efficiency hold, a reduction in energy use by up to 100x may decrease the operational costs of maintaining local AI resources, potentially reducing institutional dependency on third-party cloud providers [16]. For the healthcare and life sciences subsegment, if agentic workflows for regulatory filings scale, research compliance management may shift from the manual preparation of documentation to the systemic auditing of agent-generated submissions [13]. This represents a change in the educational substance of compliance, moving from clerical accuracy to the validation of algorithmic provenance. For institutions in regions prioritizing "frugal AI," research infrastructure management may shift toward the procurement and maintenance of smaller, specialized model architectures rather than general-purpose frontier systems [17]. The fundamental legal accountability of the Principal Investigator (PI) for research misconduct is unlikely to change, as institutional liability remains tethered to human legal personhood.

Hypothesis: In the mid term, research infrastructure management for high-compute institutions may shift toward energy-efficient local autonomy, while compliance in specialized subsegments will move from document generation to algorithmic auditing.

Far (2046–2051)

For the subsegments of life sciences and high-compute research, the scope of research compliance management may shift toward fully automated, real-time monitoring of laboratory outputs and compute patterns. However, the emergence of quantum technologies may disrupt existing AI-based administration tools, creating a period of instability in how research infrastructure is managed [18]. One significant uncertainty remains whether international standards for "model theft" and distillation will harmonize or lead to frag-

mented, regional "walled gardens" of research infrastructure. The requirement for a designated human officer to sign off on high-risk research certifications is likely to resist change due to the necessity of legal liability in state-funded research.

Hypothesis A: In the far term, research administration evolves into a highly automated, real-time oversight system where compliance is baked into the hardware and software layers of the research infrastructure.

Hypothesis B (competing): In the far term, geopolitical tensions over hardware and intellectual property lead to a fragmented research administration landscape characterized by restrictive, manual vetting and highly isolated local compute clusters.

References

1. Arxiv Cs Lg (2026-04-23). "Transparent Screening for LLM Inference and Training Impacts." https://arxiv.org/abs/2604.19757
2. Ai News (2026-04-23). "NVIDIA and Google infrastructure cuts AI inference costs." https://www.artificialintelligence-news.com/news/nvidia-and-google-infrastructure-cuts-ai-inference-costs/
3. Lesswrong (2026-04-11). "Catching illicit distributed training operations during an AI pause." https://www.lesswrong.com/posts/35yyWJnXvC2ae6NKH/catching-illicit-distributed-training-operations-during-an
4. Dev To Ai (2026-04-07). "Gemma 4: A Practical Guide to Running Frontier AI on Your Own Hardware." https://dev.to/bolgercarol/gemma-4-a-practical-guide-to-running-frontier-ai-on-your-own-hardware-5h9l
5. Hacker News (2026-04-08). "MegaTrain: Full Precision Training of 100B+ Parameter LLMs on a Single GPU." https://arxiv.org/abs/2604.05091
6. Dev To Ai (2026-04-10). "How To Navigate Enterprise GPU Shortages for AI Workloads." https://dev.to/autonainews/how-to-navigate-enterprise-gpu-shortages-for-ai-workloads-4dkj
7. The Register (2026-04-23). "AI now gobbling up power and management chips for servers." https://go.theregister.com/feed/www.theregister.com/2026/04/23/ai_now_gobbling_up_power/
8. The Register (2026-04-23). "Datadog digs down into GPU efficiency as AI costs soar." https://go.theregister.com/feed/www.theregister.com/2026/04/23/datadog_digs_down_into_gpu/
9. Ai Now Institute (2026-04-01). "North Star Data Center Policy Toolkit: State and Local Policy Interventions to Stop Rampant AI Data Center Expansion." https://ainowinstitute.org/publications/data-center-policy-guide
10. EDUCAUSE Review (2026-04-21). "White House Unveils National AI Legislative Framework." https://er.educause.edu/articles/2026/4/white-house-unveils-national-ai-legislative-framework
11. Scmp Tech (2026-04-22). "US visa system 'rolls out red carpet' for Chinese spies, congressional hearing told." https://www.scmp.com/news/china/diplomacy/article/3351062/us-visa-system-rolls-out-red-carpet-chinese-spies-congressional-hearing-told?utm_source=rss_feed
12. BBC News (2026-04-23). "White House memo claims mass AI theft by Chinese firms." https://www.bbc.com/news/articles/cpqxgxx9nrqo?at_medium=RSS&at_campaign=rss
13. Aws Ml (2026-04-08). "Human-in-the-loop constructs for agentic workflows in healthcare and life sciences." https://aws.amazon.com/blogs/machine-learning/human-in-the-loop-constructs-for-agentic-workflows-in-healthcare-and-life-sciences/
14. The Conversation Tech (2026-04-09). "AI can design and run thousands of lab experiments without human hands. Humanity isn't ready for the new risks this brings to biology." https://theconversation.com/ai-can-design-and-run-thousands-of-lab-experiments-without-human-hands-humanity-isnt-ready-for-the-new-risks-this-brings-to-biology-279191
15. Fast Company (2026-04-13). "AI is rewriting the rules of biological experiments, but safety regulations aren't keeping up." https://www.fastcompany.com/91524571/ai-biology-lab-experiments-risks?partner=rss&utm_source=rss&utm_medium=feed&utm_campaign=rss+fastcompany&utm_content=rss
16. Science Daily Comp (2026-04-05). "AI breakthrough cuts energy use by 100x while boosting accu-

racy." https://www.sciencedaily.com/releases/2026/04/260405003952.htm
17. Rest Of World (2026-04-02). "Nations priced out of Big AI are building with frugal models." https://restofworld.org/2026/frugal-ai-big-tech/?utm_source=rss&utm_medium=rss&utm_campaign=feeds
18. Scmp Tech (2026-04-14). "Chinese team shows quantum tech can disrupt AI in a real world task." https://www.scmp.com/news/china/science/article/3349995/chinese-team-shows-quantum-tech-can-disrupt-ai-real-world-task?utm_source=rss_feed

Research Delivery

capability_id: research-delivery · created: 2026-05-13 · revised: 2026-05-13 · refs: 27

Near (2026–2031)

For a subset of institutions focusing on computational, biomedical, mathematical, and astronomical research, and in specific regions including Japan and the UK, AI is altering both the delivery mechanisms and the educational substance of research production, dataset management, and output management. In research production for the life sciences, delivery is shifting toward autonomous execution via agentic discovery, where AI designs and runs lab experiments without human intervention [19], accelerates drug discovery through petabyte-scale data scanning [8], and employs agentic AI for evidence-based medical discovery [32] using specialized bioinformatics skill-sets [21] and standardized benchmarks like LABBench2 [28]. Technical research production is further accelerated by agentic models capable of executing the full stack of coding and data analysis without step-by-step human supervision [38], alongside the automation of GPU kernel optimization [4], full-precision training of large models on single GPUs [13], and fast interpretability tooling for trillion-parameter models [37]. Research output management delivery is shifting toward automated manuscript synthesis via multi-agent frameworks [18][20], AI-integrated co-authoring tools [36], and custom extraction pipelines for literature reviews [7][23]. Research dataset management is diverging regionally—where Japan's relaxed privacy laws accelerate data availability [10] while the UK faces systemic hurdles in centralizing public data [11]—and technically, through the use of privacy-preserving federated learning [24][25], the generation of controllable synthetic datasets [30], the adoption of end-to-end lineage tools [31], and cross-cloud lakehouses [34]. However, research production in behavioral and cognitive sciences faces a substance risk where LLM behavioral simulators lack causal validity [3], a failure mode compounded by "cognitive surrender" where researchers abandon logical verification in favor of AI-generated results [1]. The institutional peer-review process for research output management is unlikely to change yet, as it remains the primary mechanism for professional accountability and social legitimacy.

Hypothesis: In the near term, research delivery within AI-reliant subsegments will experience a tension between accelerated technical output and a decline in human-led logical and causal verification.

Mid (2036–2041)

For these same subsegments and regions, if current trajectories hold, research production may shift toward a verification-centric model. Because "cognitive surrender" [1] and causally invalid simulations [3] erode the validity of automated results, the substance of research production may shift from primary discovery to the rigorous verification of AI-generated hypotheses. The integration of documented human-in-the-loop validation steps, particularly in healthcare and life sciences [16], could become a requirement within research output management to mitigate biological risks [19][27] and maintain scientific legitimacy. In technical research production, the emergence of AI-to-AI communication protocols [14] may allow autonomous agents to manage the iterative loop of hypothesis testing and refinement independently. Despite these shifts, the fundamental requirement for human authors to assume legal and ethical responsibility for research outputs will likely persist due to existing institutional governance and liability frameworks.

Hypothesis: In the mid term, the primary human role in research delivery for AI-intensive fields will shift from production to a supervisory verification function to counteract the erosion of causal rigor.

Far (2046–2051)

For the computational and biomedical subsegments identified in the near term, the capability of research delivery may diverge based on the resolution of the causal validity crisis. One outcome involves the total decoupling of "production" from "authorship," where agentic loops manage the entire lifecycle of research production and dataset management, leaving humans only as ethical and legal signatories. A rival outcome involves a structural "return to rigor" where high-stakes research production mandates the removal of agentic autonomy in the hypothesis-generation phase to restore causal certainty. The impact of quantum technology on AI-driven research production remains an unresolvable uncertainty that could

either accelerate these trajectories or render current agentic frameworks obsolete [29]. Even at this horizon, the social requirement for human-attributable "discovery" is likely to resist change, as the prestige economy of higher education depends on individual intellectual contribution.

Hypothesis A: In the far term, research delivery becomes a fully autonomous agentic process where the human role is reduced to a legal and ethical audit function.

Hypothesis B (competing): In the far term, a systemic failure in AI-generated causal validity forces a return to human-led research production for all high-stakes scientific discovery.

References

1. Ars Technica (2026-04-03). ""Cognitive surrender" leads AI users to abandon logical thinking, research finds." https://arstechnica.com/ai/2026/04/research-finds-ai-users-scarily-willing-to-surrender-their-cognition-to-llms/
2. Arxiv Cs Cy (2026-04-06). "When simulations look right but causal effects go wrong: Large language models as behavioral simulators." https://arxiv.org/abs/2604.02458
3. Marktechpost (2026-04-06). "RightNow AI Releases AutoKernel: An Open-Source Framework that Applies an Autonomous Agent Loop to GPU Kernel Optimization for Arbitrary PyTorch Models." https://www.marktechpost.com/2026/04/06/rightnow-ai-releases-autokernel-an-open-source-framework-that-applies-an-autonomous-agent-loop-to-gpu-kernel-optimization-for-arbitrary-pytorch-models/
4. Dev To Ai (2026-04-07). "Automate Your Literature Review: Build a Custom AI Extraction Pipeline." https://dev.to/ken_deng_ai/automate-your-literature-review-build-a-custom-ai-extraction-pipeline-1dmd
5. The Conversation Tech (2026-04-07). "AI is reengineering drug discovery by speeding up testing and scanning petabytes of data for connections between diseases." https://theconversation.com/ai-is-reengineering-drug-discovery-by-speeding-up-testing-and-scanning-petabytes-of-data-for-connections-between-diseases-274693
6. The Register (2026-04-08). "Japan relaxes privacy laws to make itself the 'easiest country to develop AI'." https://go.theregister.com/feed/www.theregister.com/2026/04/08/japan_privacy_law_changes_ai/
7. The Register (2026-04-08). "UK's grand plan to fuel AI with public data faces uphill battle." https://go.theregister.com/feed/www.theregister.com/2026/04/08/national_data_library_plan/
8. Hacker News (2026-04-08). "MegaTrain: Full Precision Training of 100B+ Parameter LLMs on a Single GPU." https://arxiv.org/abs/2604.05091
9. The Register (2026-04-08). "Talk ain't cheap: DARPA offers grants for new AI-to-AI communication protocol." https://go.theregister.com/feed/www.theregister.com/2026/04/08/darpa_wants_ai_agent_communication/
10. Aws Ml (2026-04-08). "Human-in-the-loop constructs for agentic workflows in healthcare and life sciences." https://aws.amazon.com/blogs/machine-learning/human-in-the-loop-constructs-for-agentic-workflows-in-healthcare-and-life-sciences/
11. Marktechpost (2026-04-09). "Google AI Research Introduces PaperOrchestra: A Multi-Agent Framework for Automated AI Research Paper Writing." https://www.marktechpost.com/2026/04/08/google-ai-research-introduces-paperorchestra-a-multi-agent-framework-for-automated-ai-research-paper-writing/
12. The Conversation Tech (2026-04-09). "AI can design and run thousands of lab experiments without human hands. Humanity isn't ready for the new risks this brings to biology." https://theconversation.com/ai-can-design-and-run-thousands-of-lab-experiments-without-human-hands-humanity-isnt-ready-for-the-new-risks-this-brings-to-biology-279191
13. Dev To Ai (2026-04-09). "Building a Multi-Agent Research System with LangGraph: How I Taught Three AI Agents to Collaborate." https://dev.to/benturtle/building-a-multi-agent-research-system-with-langgraph-how-i-taught-three-ai-agents-to-collaborate-2ng6

14. Dev To Ai (2026-04-10). "Add 197 Bioinformatics Skills to Claude Code with SciAgent-Skills." https://dev.to/gentic_news/add-197-bioinformatics-skills-to-claude-code-with-sciagent-skills-32gh
15. OpenAI (2026-04-10). "Research with ChatGPT." https://openai.com/academy/search-and-deep-research
16. Dev To Ai (2026-04-11). "Federated Learning in 2026: Privacy-Preserving AI." https://dev.to/lufumeiying/federated-learning-in-2026-privacy-preserving-ai-5f9j
17. Dev To Ai (2026-04-11). "Why N=1 Rare Disease Sites Are Excluded from Federated Learning by Architecture — and What Outcome Routing Does Instead." https://dev.to/roryqis/why-n1-rare-disease-sites-are-excluded-from-federated-learning-by-architecture-and-what-outcome-b9c
18. Fast Company (2026-04-13). "AI is rewriting the rules of biological experiments, but safety regulations aren't keeping up." https://www.fastcompany.com/91524571/ai-biology-lab-experiments-risks?partner=rss&utm_source=rss&utm_medium=feed&utm_campaign=rss+fastcompany&utm_content=rss
19. Arxiv Cs Ai (2026-04-14). "LABBench2: An Improved Benchmark for AI Systems Performing Biology Research." https://arxiv.org/abs/2604.09554
20. Scmp Tech (2026-04-14). "Chinese team shows quantum tech can disrupt AI in a real world task." https://www.scmp.com/news/china/science/article/3349995/chinese-team-shows-quantum-tech-can-disrupt-ai-real-world-task?utm_source=rss_feed
21. Marktechpost (2026-04-21). "Google Introduces Simula: A Reasoning-First Framework for Generating Controllable, Scalable Synthetic Datasets Across Specialized AI Domains." https://www.marktechpost.com/2026/04/21/google-introduces-simula-a-reasoning-first-framework-for-generating-controllable-scalable-synthetic-datasets-across-specialized-ai-domains/
22. Aws Ml (2026-04-21). "End-to-end lineage with DVC and Amazon SageMaker AI MLflow apps." https://aws.amazon.com/blogs/machine-learning/end-to-end-lineage-with-dvc-and-amazon-sagemaker-ai-mlflow-apps/
23. Arxiv Cs Ai (2026-04-21). "DeepER-Med: Advancing Deep Evidence-Based Research in Medicine Through Agentic AI." https://arxiv.org/abs/2604.15456
24. Dev To Ai (2026-04-23). "The Announcement Everyone Slept On at Google Cloud Next '26: The Cross-Cloud Lakehouse." https://dev.to/precious_pendo_911264e7ce/the-announcement-everyone-slept-on-at-google-cloud-next-26-the-cross-cloud-lakehouse-233n
25. Nvidia Blog (2026-04-23). "Making Sense of the Early Universe." https://blogs.nvidia.com/blog/ai-gpu-early-universe-astronomy/
26. The Register (2026-04-23). "Microsoft gives your Word documents an AI co-author you didn't ask for." https://go.theregister.com/feed/www.theregister.com/2026/04/23/microsoft_gives_your_word_documents/
27. Lesswrong (2026-04-23). "vLLM-Lens: Fast Interpretability Tooling That Scales to Trillion-Parameter Models." https://www.lesswrong.com/posts/3bs27nZQuEcKhXf7q/vllm-lens-fast-interpretability-tooling-that-scales-to
28. Marktechpost (2026-04-23). "OpenAI Releases GPT-5.5, a Fully Retrained Agentic Model That Scores 82.7% on Terminal-Bench 2.0 and 84.9% on GDPval." https://www.marktechpost.com/2026/04/23/openai-releases-gpt-5-5-a-fully-retrained-agentic-model-that-scores-82-7-on-terminal-bench-2-0-and-84-9-on-gdpval/

Research Funding

capability_id: research-funding · created: 2026-05-02 · revised: 2026-05-02 · refs: 2

Near (2026–2031)

For US federal research funding and specialized grant-seeking subsegments, AI is being applied to both the management of existing funds and the identification of new leads. In the US federal context, the process of research funds management may be subject to automated thematic filtering where large language models are used to identify grants for elimination based on ideological or thematic markers [1]. This represents a change in educational substance, as the criteria for funding validity shift from expert peer review toward algorithmic keyword alignment. Simultaneously, for a subsegment of grant seekers, the process of research funding identification is shifting from manual directory searches to AI-automated prospecting and personalized lead generation [2]. This change is primarily one of delivery mechanism, replacing manual research with strategic, automated pipeline building. The legal mechanisms for the actual disbursement of funds are unlikely to change yet due to the rigidity of federal appropriation laws and statutory requirements.

Hypothesis: In the near term, US federal research funds management may incorporate automated filtering for grant elimination, while grant-seeking subsegments may adopt AI-driven automation for research funding identification.

Mid (2036–2041)

For US federal research funding and specialized grant-seeking subsegments, if current trajectories hold, the interaction between automated identification and automated filtering may create a feedback loop of prompt optimization. Extrapolating from current use, researchers seeking federal funds may treat research funding identification as a prompt-engineering exercise, tailoring proposal language to avoid the algorithmic triggers used in research funds management [1]. In the grant-seeking subsegment, the use of automated prospecting may lead to a saturation of high-probability leads, potentially increasing the volume of personalized but low-substance outreach to foundations [2]. Despite these shifts, the administrative reporting requirements for how funds are spent will likely persist unchanged because they serve as the primary mechanism for legal and financial accountability.

Hypothesis: In the mid term, research funding identification may shift toward a prompt-optimization model where proposals are engineered to bypass algorithmic filters used in fund management.

Far (2046–2051)

For US federal research funding and specialized grant-seeking subsegments, the long-term interaction between automated prospecting and algorithmic elimination remains uncertain. One possibility is that these tools evolve into sophisticated systems that optimize grant portfolios based on complex data patterns, although it is unknown if these patterns correlate with genuine scientific value. A rival outcome involves the systemic erosion of the social legitimacy of federal grants if the management process is perceived as purely ideological rather than meritocratic. The fundamental requirement for financial auditing and the prevention of fraud will likely resist change even at this horizon, as these functions are tied to legal liability rather than thematic preference.

Hypothesis A: In the far term, US federal research funding management may transition to a bounded autonomous system that optimizes grant portfolios based on predefined state objectives.

Hypothesis B (competing): In the far term, the perceived lack of meritocratic rigor in algorithmic filtering may lead to a migration of research talent toward non-governmental funding sources.

References

1. Fast Company (2026-04-14). "What good AI in government actually looks like." https://www.fastcompany.com/91525531/grantwell-what-good-ai-in-government-actually-looks-like?partner=rss&u

tm_source=rss&utm_medium=feed&utm_campaign=rss+fastcompany&utm_content=rss
2. Dev To Ai (2026-04-22). "From Scattershot to Sniper: AI for Smarter Grant Lead Generation." https://dev.to/ken_deng_ai/from-scattershot-to-sniper-ai-for-smarter-grant-lead-generation-4if4

Research Impact

capability_id: research-impact · created: 2026-04-11 · revised: 2026-04-11 · refs: 2

Near (2026–2031)

For researchers managing non-commercial research impact via digital dissemination, the process of visibility management may shift from search engine optimization (SEO) toward generative engine optimization (GEO) [1]. Because users increasingly employ LLMs as a discovery layer that synthesizes answers from a retrieval set, the mechanism for achieving impact shifts from winning a click to becoming a verifiable and extractable source for model synthesis [1]. Evidence suggests that legacy web authority metrics, such as Domain Authority (DA), no longer reliably predict whether AI platforms will cite a specific source [2]. This decoupling of domain prestige from AI citability suggests that non-commercial research impact management must prioritize content structure over traditional domain authority [2]. This represents a change in the delivery mechanism of the capability, as researchers may alter how they structure digital content to ensure it is "RAG-friendly" [1]. However, core metrics for research impact, such as h-index or formal citations in peer-reviewed journals, are unlikely to change in this period due to their deep embedding in institutional tenure and funding regulations.

Hypothesis: In the near term, researchers utilizing digital dissemination may adapt their visibility strategies to prioritize machine-readability and structural verifiability over traditional domain authority to increase their presence within generative answer engines.

Mid (2036–2041)

For the same subset of researchers focused on web-based dissemination, if current trajectories hold, the process of non-commercial research impact management could transition toward formalized GEO standards. Extrapolating from the requirement for verifiable sources in RAG-based systems [1], researchers may adopt specific metadata schemas to ensure their work is correctly attributed during LLM synthesis. This could shift the educational substance of "impact" from simple reach (clicks or views) to "synthesized utility," where the value of research is measured by its role as a foundational component of AI-generated answers. Despite this, the process of commercial research outcomes—such as patent filing and licensing—will likely persist unchanged because these rely on legal protections and proprietary disclosures rather than public discovery layers.

Hypothesis: In the mid term, non-commercial research impact for web-active researchers may be increasingly evaluated based on the frequency and accuracy of their work's synthesis within generative engines.

Far (2046–2051)

For researchers using digital discovery layers, the capability of managing non-commercial research impact remains subject to the tension between algorithmic visibility and institutional trust. One plausible outcome is that the retrieval set becomes the primary arbiter of academic influence, effectively automating the discovery process for a large portion of non-commercial impact. Conversely, a rival outcome may emerge if the over-optimization of research for AI synthesis leads to a counterproductive erosion of trust, where the "gaming" of GEO undermines the perceived validity of the research itself. It remains unknown whether funding bodies will eventually replace traditional citation counts with generative synthesis metrics. The requirement for human-peer-validated novelty will likely resist change even at this horizon, as it provides the social legitimacy that algorithmic synthesis cannot generate.

Hypothesis A: In the far term, digital research impact may be primarily mediated through the precision and frequency with which a scholar's work is retrieved and synthesized by autonomous agents.

Hypothesis B (competing): In the far term, institutional distrust in AI-mediated visibility may lead to a re-valuation of gated, non-algorithmic verification processes to ensure research impact is not artificially inflated via GEO.

References

1. Dzone Ai (2026-04-06). "Engineering for the Answer Engine: GEO for RAG-Friendly Web Apps (TalentHacked.com Case Study)." https://feeds.dzone.com/link/23558/17314715/engineering-for-the-answer-engine-geo-for-rag-frie
2. Dev To Ai (2026-04-10). "Why DA Is Irrelevant for AI Citations (Data from 7 Site Audits)." https://dev.to/chudi_nnorukam/why-da-is-irrelevant-for-ai-citations-data-from-7-site-audits-1il0

Research Improvement

capability_id: research-improvement · created: 2026-04-23 · revised: 2026-04-23 · refs: 3

Near (2026–2031)

For a subset of research quality management involving biology research, behavioral simulations, and the scientific drafting process, the introduction of specialized benchmarks and automated feedback mechanisms is altering validation protocols. In biology research, the deployment of standardized benchmarks such as LABBench2 enables more precise evaluation of AI systems, shifting the delivery mechanism of quality management toward automated performance auditing [2]. In behavioral simulation, a disconnect between the visual plausibility of LLM outputs and their actual causal accuracy necessitates new validation criteria that prioritize causal verification over output coherence [1]. This represents a change in research substance, as "valid evidence" in simulations must shift from plausibility to verified causal effects [1]. Additionally, the use of LLMs to generate constructive feedback on scientific papers based on author responses is altering the research quality management process during the drafting and revision phase [3]. Researcher performance management is unlikely to change in this horizon, as institutional metrics for productivity typically rely on output volume rather than the technical fidelity of AI-assisted methodologies.

Hypothesis: In the near term, research quality management for specific AI-driven subsegments may shift toward requiring standardized benchmark validation, explicit causal verification, and AI-mediated constructive feedback.

Mid (2036–2041)

For the same subsegments of biology, behavioral, and drafting processes, if current trajectories toward standardized benchmarking and causal auditing hold, research quality management may integrate these technical audits as mandatory prerequisites for publication [1][2]. Extrapolating from the use of tools like LABBench2 and GoodPoint, the process of research quality management could transition from human-led peer review of final results to a technical verification of the AI's operational logic, causal architecture, and feedback iterations [1][2][3]. This shift moves the point of intervention in quality management from the end of the research cycle to the tool-selection and execution phases. The reliance on institutional review boards for final methodological approval is likely to persist, as these bodies provide the social legitimacy required for high-stakes scientific claims.

Hypothesis: In the mid term, the validation of AI-driven research in these subsegments may evolve into a formal technical auditing process that separates surface-level plausibility from structural validity.

Far (2046–2051)

For researchers utilizing synthetic behavioral agents and AI-driven biological discovery, the stability of research quality management depends on whether the gap between plausibility and causal accuracy is fundamentally resolved in model architectures [1]. One plausible outcome involves a specialized regime of "synthetic provenance," where AI-generated data is strictly tiered by its verified causal reliability according to evolving benchmarks [2]. A rival outcome suggests a systemic devaluation of LLM-simulated data for high-stakes causal claims if persistent inaccuracies remain, leading to a renewed institutional preference for empirical human data. It remains unknown whether future architectures can inherently solve the plausibility-accuracy gap or if external, independent verification will always be a requirement. The fundamental requirement for empirical grounding in behavioral and biological sciences will resist change even at this horizon.

Hypothesis A: In the far term, research quality management may successfully incorporate a tiered system of synthetic data reliability based on rigorous, automated causal verification.

Hypothesis B (competing): In the far term, a persistent gap between plausibility and causal accuracy may lead to the systemic devaluation of AI-simulated data in favor of empirical human data.

References

1. Arxiv Cs Cy (2026-04-06). "When simulations look right but causal effects go wrong: Large language models as behavioral simulators." https://arxiv.org/abs/2604.02458
2. Arxiv Cs Ai (2026-04-14). "LABBench2: An Improved Benchmark for AI Systems Performing Biology Research." https://arxiv.org/abs/2604.09554
3. Arxiv Cs Ai (2026-04-16). "GoodPoint: Learning Constructive Scientific Paper Feedback from Author Responses." https://arxiv.org/abs/2604.11924

Research Opportunity Analysis

capability_id: research-opportunity-analysis · created: 2026-04-10 · revised: 2026-04-10 · refs: 1

Near (2026–2031)

For research-intensive environments utilizing agentic AI tools, the process of research prioritisation may shift toward automated gap detection. The use of parallel AI research agents allows for the systematic mapping of knowledge gaps across vast scientific corpora [1]. This represents a change in the delivery mechanism of research prioritisation, replacing some manual literature synthesis with agentic scanning. Because this process identifies informational voids rather than strategic value, the final selection of research directions will likely remain a human-led activity governed by resource constraints and institutional priorities. The social legitimacy of a research agenda continues to depend on expert human validation, making a fully automated prioritisation process unlikely in this horizon.

Hypothesis: AI agents will support the initial stage of knowledge gap identification within research prioritisation in the near term.

Mid (2036–2041)

For institutions implementing agentic workflows in their research offices, if current trajectories in automated mapping hold, these tools could alter the process of research programme development. Extrapolating from current demonstrations of parallel agentic gap analysis [1], the identification of "white space" in scientific literature may become a standardized input for designing multi-year research programmes. This shift may move the substance of programme development away from intuition-based hypothesis generation toward a more data-driven model of novelty detection. However, the requirement for peer-reviewed validation of a research gap's significance will likely persist, as AI cannot yet simulate the social consensus required for academic legitimacy.

Hypothesis: Research prioritisation may evolve into a hybrid model where AI agents propose potential knowledge gaps for human curation in the mid term.

Far (2046–2051)

For research-intensive environments, the long-term impact on research opportunity analysis remains uncertain, specifically regarding whether AI can identify conceptual paradigm shifts or merely informational gaps. One possibility is that automated gap mapping becomes the primary driver of research agendas, potentially leading to a counterproductive cycle where researchers pursue a high volume of superficially novel but low-impact "gaps" identified by AI. Conversely, a reaction against this automation may occur, where the value of "theory-driven" rather than "gap-driven" research is re-emphasized to maintain intellectual substance. The physical constraint of empirical validation—the need to conduct actual experiments or observations—will resist automation even at this horizon.

Hypothesis A: AI-driven gap analysis becomes the primary mechanism for research prioritisation in the far term.

Hypothesis B (competing): The proliferation of AI-identified gaps leads to a crisis of significance, triggering a return to human-centric, theory-led research prioritisation in the far term.

References

1. Dev To Ai (2026-04-10). "We mapped 153 gaps in science using 5 parallel AI research agents." https://dev.to/minimal-architect/we-mapped-153-gaps-in-science-using-5-parallel-ai-research-agents-5fdk

Research Publications

capability_id: research-publications · created: 2026-04-11 · revised: 2026-04-11 · refs: 7

Near (2026–2031)

For research publications broadly, research publication management is increasingly defined by a tension between automated drafting and the deployment of provenance-tracking infrastructure. Publishers are employing AI-detection tools to maintain the integrity of submitted research outputs [1], while the emergence of AI citation registries suggests a shift toward using structured information infrastructure to verify the authenticity of references [2]. For a subsegment of researchers, the delivery mechanism of research publication management is shifting toward the use of multi-agent frameworks, such as PaperOrchestra, to automate the drafting of manuscripts [5]. Regarding research output reporting, a subsegment of researchers is adapting output formatting—utilizing direct Q&A and tables—to increase visibility within AI search engines [6]. This shift is supported by the development of AI visibility infrastructure designed to interface directly with AI systems [4] and evidence that AI citations are driven by specific "citability" factors rather than traditional domain authority (DA) [7]. In institution-specific settings, such as the Alignment Journal, experimental policies are introducing reviewer compensation to modify the peer-review workflow [3]. The social legitimacy of the human-led peer-review process is unlikely to change in this window because institutional tenure and promotion regimes rely on human intellectual accountability.

Hypothesis: In the near term, research publication management will be characterized by a technical arms race between multi-agent drafting tools and the infrastructure required to detect and verify AI-generated content.

Mid (2036–2041)

For research publications broadly, if current trajectories in registry adoption and detection instability hold, the mechanism of verification may shift from reactive detection to proactive provenance. Extrapolating from the limitations of binary AI-text detection [1], publishers may require research output reporting to be anchored in verified citation registries [2] to prevent the proliferation of fabricated evidence. Research output reporting may further diverge from the static PDF toward structured entity formats to maintain discoverability in AI-driven knowledge graphs [4][6]. The journal article is likely to persist as the primary unit of currency for academic promotion due to high institutional inertia in tenure committees. However, the delivery mechanism of the peer-review process may integrate structured transparency logs to counteract the opacity of AI-assisted drafting [5].

Hypothesis: In the mid term, the focus of research publication management may move from the detection of synthetic text to the systematic verification of authorship and citation provenance through structured registries.

Far (2046–2051)

For research publications broadly, the long-term viability of the written manuscript as the primary research output remains unknown if the gap between automated generation [5] and detection [1] closes completely. One plausible outcome is a transition where the primary research output becomes a verifiable computational model or a registry-linked data stream [2], with written text serving only as a secondary, AI-generated summary for human consumption [4]. A rival outcome involves the erosion of trust in traditional publication management, leading to a fragmentation of the reward system as traditional journals fail to distinguish human discovery from synthetic synthesis. It remains unknown whether visibility infrastructure [4] can scale sufficiently to maintain a global, trusted record of discovery. The fundamental requirement for a verifiable and immutable record of intellectual priority resists change even at this horizon.

Hypothesis A: In the far term, the primary research output may shift from the written manuscript to verifiable computational models or data streams linked to citation registries.

Hypothesis B (competing): In the far term, the inability to distinguish human from synthetic discovery

may lead to a collapse of the traditional journal system and a fragmented, localized reward structure for research.

References

1. Dev To Ai (2026-04-06). "Will AI-Written Papers Be Detected? The Truth About AIGC Detection in 2026." https://dev.to/yu_z_ed6c1c576593c5e7d41f/will-ai-written-papers-be-detected-the-truth-about-aigc-detection-in-2026-5b50
2. Dev To Ai (2026-04-06). "AI Citation Registries as Information Infrastructure for AI Systems." https://dev.to/aigistry/ai-citation-registries-as-information-infrastructure-for-ai-systems-50o6
3. Lesswrong (2026-04-07). "An Alignment Journal: Features and policies." https://www.lesswrong.com/posts/hvq7amw8FKZxEeKqA/an-alignment-journal-features-and-policies
4. Dev To Ai (2026-04-08). "Building AI Visibility Infrastructure: Inside Jonomor's Architecture." https://dev.to/jonomor_ecosystem/building-ai-visibility-infrastructure-inside-jonomors-architecture-3enp
5. Marktechpost (2026-04-09). "Google AI Research Introduces PaperOrchestra: A Multi-Agent Framework for Automated AI Research Paper Writing." https://www.marktechpost.com/2026/04/08/google-ai-research-introduces-paperorchestra-a-multi-agent-framework-for-automated-ai-research-paper-writing/
6. Dev To Ai (2026-04-10). "How to Get Cited by AI Search Engines: ChatGPT, Perplexity & Claude in 2026." https://dev.to/iris1031/how-to-get-cited-by-ai-search-engines-chatgpt-perplexity-claude-in-2026-21p0
7. Dev To Ai (2026-04-10). "Why DA Is Irrelevant for AI Citations (Data from 7 Site Audits)." https://dev.to/chudi_nnorukam/why-da-is-irrelevant-for-ai-citations-data-from-7-site-audits-1il0

Research Training

capability_id: research-training · created: 2026-04-23 · revised: 2026-04-23 · refs: 6

Near (2026–2031)

For research training broadly, the researcher training and development process may integrate AI-driven search and deep research tools to automate the identification of sources and the generation of structured literature syntheses [3][4]. This modification affects the delivery mechanism of the literature review process rather than its core intellectual requirement. For a subsegment of researchers engaged in quantitative analysis, postgraduate research student skills development may integrate AI-driven data analysis tools to automate dataset exploration and initial visualization [5]. Within the specific subsegment of machine learning research, student skills development may shift from a primary focus on theoretical mathematical derivation toward compute-intensive engineering and the management of large-scale empirical experiments [1]. This represents a change in educational substance, as progress in the field increasingly relies on compute scaling over mathematically principled architectures [1]. In the subsegment of scientific writing and peer review training, the development of researcher skills may incorporate AI systems designed to generate constructive feedback based on author responses [6]. The institutional requirement for a written thesis and a formal oral defense is unlikely to change in the near term due to the high inertia of degree-awarding regulations and accreditation standards.

Hypothesis: In the near term, general research training will integrate AI-driven sourcing and synthesis, while technical subsegments will pivot toward compute engineering, automated data analysis, and AI-mediated feedback loops.

Mid (2036–2041)

For the same broad and technical cohorts, the postgraduate research student skills development process may transition from a focus on information retrieval to the rigorous validation of AI-curated syntheses if current trajectories hold [3][4]. Extrapolating from the adoption of agentic terminal tools, researcher training in technical subsegments may replace manual codebase maintenance with the orchestration of autonomous coding agents as a primary technical competency [2]. This shift may alter the substance of postgraduate research supervisor development, as supervisors may move from guiding manual methodology to validating agent-generated empirical workflows. The necessity for the researcher to frame a novel research question and maintain an original theoretical contribution is likely to persist unchanged.

Hypothesis: In the mid term, research training will shift toward the validation of automated synthesis and the orchestration of agentic systems, shifting supervisory focus toward workflow validation.

Far (2046–2051)

For the machine learning and technical research subsegments, the long-term trajectory of skills development depends on whether compute-intensive scaling continues to yield returns or reaches a threshold of diminishing returns [1]. If scaling remains the primary driver of discovery, research training may evolve into a specialized discipline of high-performance infrastructure management and agent orchestration. Conversely, if scaling returns diminish, a rival trajectory may see a renewed emphasis on mathematically principled architectures to enable new breakthroughs [1]. It remains unknown whether the human role in general researcher training will shift from primary synthesizer to an overseer of automated discovery systems [3][4]. The institutional requirement for peer-reviewed validation of research results will likely resist change even at this horizon.

Hypothesis A: In the far term, research training in technical subsegments may fully transition into a specialized form of high-performance infrastructure and agentic system management.

Hypothesis B (competing): In the far term, a decline in the efficacy of compute-scaling may force a return to mathematically grounded researcher training to achieve new architectural breakthroughs.

References

1. The Gradient (2024). "Shape, Symmetries, and Structure: The Changing Role of Mathematics in Machine Learning Research." https://thegradient.pub/shape-symmetry-structure/
2. Dev To Ai (2026). "Claude Code Cheatsheet: The Terminal AI Agent That Actually Codes." https://dev.to/mibii/claude-code-cheatsheet-the-terminal-ai-agent-that-actually-codes-36bk
3. OpenAI (2026). "ChatGPT for research." https://openai.com/academy/research
4. OpenAI (2026). "Research with ChatGPT." https://openai.com/academy/search-and-deep-research
5. OpenAI (2026). "Analyzing data with ChatGPT." https://openai.com/academy/data-analysis
6. Arxiv Cs Ai (2026). "GoodPoint: Learning Constructive Scientific Paper Feedback from Author Responses." https://arxiv.org/abs/2604.11924

Strategy Management

capability_id: strategy-management · created: 2026-05-09 · revised: 2026-05-09 · refs: 8

Near (2026–2031)

For institutions in China, state-funded research bodies, public universities subject to state-mandated reviews, adopters of AI-driven synthesis tools, and institutions broadly redefining performance metrics, strategy management is shifting toward external alignment and tool-assisted synthesis. In China, state targets for AI deployment are being integrated into strategic plan development and strategic plan management via national five-year planning cycles [1]. For state-funded research institutes, as demonstrated by the Alan Turing Institute, funding bodies may mandate revisions to strategic plan development to ensure adherence to legal duties and value-for-money metrics [2]. In regions utilizing labor automation data, the availability of forecasting hubs may alter business horizon scanning by providing structured data on workforce shifts to inform long-term planning [7]. For public institutions subject to state-mandated viability reviews, as seen at Iowa State University, corporate performance management is shifting toward the use of enrollment data to drive program closures or mergers [6]. For institutions adopting AI-driven consulting platforms, the processes of strategic reporting and business horizon scanning may be automated through tools that synthesize competitive intelligence into high-level analysis [3][4]. Broadly, the use of generative AI for deriving market insights may accelerate business horizon scanning by reducing the time lag between detecting market shifts and initiating a strategic response [5], while a shift in corporate performance management is emerging that seeks to replace routine cognitive KPIs with metrics reflecting human-centric value [8]. These changes represent a mix of delivery-mechanism changes—such as the use of AI for reporting and horizon scanning—and changes to educational substance, specifically in how corporate performance management prioritizes enrollment viability [6] and redefined value metrics [8]. The fundamental legal fiduciary duties of governing boards and the structural cadence of multi-year national planning cycles are unlikely to change in this period.

Hypothesis: In the near term, strategy management in these settings may be increasingly defined by external mandates and the use of automated synthesis to accelerate market intelligence, reporting, and the redefinition of performance metrics.

Mid (2036–2041)

For institutions in China, state-funded research bodies, public universities under state oversight, and adopters of AI strategy tools, if current trajectories hold, corporate performance management may shift toward the quantitative measurement of external benchmarks and AI-generated KPIs. Extrapolating from state and funder mandates [1][2] and mandated program reviews [6], strategic reporting could move from an internal reflection to a compliance verification tool, altering the educational substance of how institutional success is defined. For institutions using automated strategy platforms, the process of strategic plan development may shift from an act of primary synthesis to an act of auditing AI-generated outputs [3][4][5]. If the shift toward meaning-based KPIs persists [8], strategic vision development may diverge between institutions that prioritize AI-driven efficiency and those that prioritize human-centric value creation. The requirement for institutions to maintain a distinct institutional identity within their strategic vision development is likely to persist, though it will operate within boundaries set by external authorities and tool constraints.

Hypothesis: In the mid term, corporate performance management in these settings may shift toward the quantitative verification of external benchmarks and the auditing of AI-generated strategic outputs.

Far (2046–2051)

For institutions in China, state-funded research bodies, public universities under state oversight, and early adopters of automated strategy platforms, the long-term state of strategy management depends on the balance between external steering and institutional autonomy. In state-steered environments, strategic plan management may become a real-time alignment process where institutional goals are dynamically

updated based on national AI deployment targets [1]. Conversely, a rival outcome may emerge where the perceived counterproductivity of automated strategy leads to a return to human-led strategic vision development as a primary source of institutional legitimacy and differentiation [8]. It remains unknown whether AI can effectively model the "meaning" and "connection" required for high-level strategic visioning, or if such capabilities remain exclusively human. The necessity for institutions to obtain social legitimacy from external stakeholders and regulatory bodies will resist change even at this horizon.

Hypothesis A: In the far term, strategy management in state-steered environments may evolve into a real-time alignment mechanism driven by national deployment targets and automated compliance.

Hypothesis B (competing): In the far term, the erosion of institutional distinctiveness caused by automated strategy may trigger a return to human-centric strategic visioning as a critical tool for social legitimacy.

References

1. Ai News (2026-04-02). "China's Five-Year Plan details the targets for AI deployment." https://www.artificialintelligence-news.com/news/chinas-five-year-plan-details-the-targets-for-ai-deployment/
2. Guardian Tech (2026-04-03). "UK's leading AI research institute told to make 'significant' changes." https://www.theguardian.com/technology/2026/apr/03/uk-alan-turing-institute-ai-research-underperforming-ukri
3. TechCrunch (2026-04-07). "AI startup Rocket offers vibe McKinsey-style reports at a fraction of the cost." https://techcrunch.com/2026/04/06/indian-startup-rocket-wants-its-ai-to-do-mckinsey-style-consulting-at-a-fraction-of-the-cost/
4. Dev To Ai (2026-04-07). "AI startup Rocket offers vibe McKinsey-style reports at a fraction of the cost." https://dev.to/minimal-architect/ai-startup-rocket-offers-vibe-mckinsey-style-reports-at-a-fraction-of-the-cost-1edd
5. Mit Sloan (2026-04-08). "Gain Consumer Insight With Generative AI." https://sloanreview.mit.edu/article/gain-consumer-insight-with-generative-ai/
6. Higher Ed Dive (2026-04-10). "Iowa State plans to cut or merge 23 programs after mandated review." https://www.highereddive.com/news/iowa-state-university-program-cuts-review-regents/817262/
7. Lesswrong (2026-04-21). "Introducing the Labor Automation Forecasting Hub." https://www.lesswrong.com/posts/So7wrd6djpQcoP4rz/introducing-the-labor-automation-forecasting-hub-1
8. Fast Company (2026-04-23). "The Age of AI means we need to throw out our old KPIs and replace them with new ones." https://www.fastcompany.com/91521422/ai-kpis-new-metrics?partner=rss&utm_source=rss&utm_medium=feed&utm_campaign=rss+fastcompany&utm_content=rss

Student Administration

capability_id: student-administration · created: 2026-04-11 · revised: 2026-04-11 · refs: 1

Near (2026–2031)

For public universities subject to state-mandated program reviews, as demonstrated by the experience at Iowa State University [1], the process of programme transfer management may face sudden surges in administrative volume. This occurs because mandated closures of low-enrollment programs require the systematic reassignment of currently enrolled students to merged or alternative programs [1]. This represents a change to the delivery mechanism of student administration—specifically the volume and speed of processing—rather than a change in the educational substance of the capability. Student records & details management will likely remain stable in its core function, though the frequency of updates to degree paths will increase. It is unlikely that the fundamental regulatory requirements for degree conferral will change yet, as these remain bound by external accreditation standards.

Hypothesis: In the near term, for a subset of public institutions facing state-mandated program cuts, programme transfer management may become a significant administrative bottleneck.

Mid (2036–2041)

For public universities facing similar regulatory pressures, if current trajectories of mandated program reviews hold, the increased load on programme transfer management could lead to the adoption of automated matching tools to suggest alternative programs for displaced students. This would further shift the delivery mechanism of student administration toward algorithmic decision-support. Extrapolating from current administrative constraints, student financial administration may require tighter integration with transfer workflows to manage the reallocation of program-specific scholarships. However, the final validation of enrolment status management will likely persist as a human-governed process to ensure legal compliance with state funding mandates.

Hypothesis: In the mid term, for institutions experiencing frequent program restructuring, the delivery mechanism of programme transfer management may shift toward automated suggestion systems.

Far (2046–2051)

For institutions subject to state-level programmatic volatility, one plausible outcome is that programme transfer management becomes a dynamic process that re-aligns student enrolments in real-time based on labor market signals and state mandates. A rival outcome is that the persistent instability of program offerings leads to a shift toward broader, degree-agnostic enrolment status management, which would reduce the frequency of specific programme transfer events. It remains unknown whether state regents will maintain the authority to mandate such cuts or if institutional autonomy over program design will be restored. The legal requirement for permanent, immutable student records will likely resist change even at this horizon.

Hypothesis A: In the far term, for public universities in volatile regulatory environments, programme transfer management may be integrated into dynamic labor-market alignment systems.

Hypothesis B (competing): In the far term, the administrative burden of frequent program cuts may drive a shift toward modular, degree-agnostic enrolment status management, reducing the need for traditional programme transfer processes.

References

1. Higher Ed Dive (2026). "Iowa State plans to cut or merge 23 programs after mandated review." https://www.highereddive.com/news/iowa-state-university-program-cuts-review-regents/817262/

Student Admission Management

capability_id: student-admission-management · created: 2026-04-25 · revised: 2026-04-25 · refs: 1

Near (2026–2031)

For a subsegment of institutions where financial aid offer letters are characterized by opacity and confusing terminology [1], the process of offer, acceptance & quota management may be impacted by a shift in how students decode these documents. Because the current delivery mechanism of financial aid communication often omits critical cost information [1], students may increasingly employ large language models to normalize and compare disparate offer formats. This represents a change in the delivery mechanism of the capability, as the point of interpretation shifts from the institution's intent to an AI-mediated analysis. However, the student application management process is unlikely to change in the near term due to the institutional inertia of established admissions cycles and regulatory reporting requirements.

Hypothesis: In the near term, the offer, acceptance & quota management process for a subsegment of institutions may see a shift in power as students use AI to decode opaque financial aid offers.

Mid (2036–2041)

Extrapolating from the current lack of transparency in offer letters [1], some institutions may attempt to use AI-driven personalization to optimize quota management by tailoring the presentation of financial aid to specific student psychographics. This shift in delivery mechanism could increase the risk of counterproductivity, where the pursuit of higher yield rates through calculated opacity undermines the social legitimacy of the admission process. Despite these potential shifts, the legal frameworks governing federal financial aid reporting will likely remain a stabilizing constraint, preventing the total removal of standard disclosures.

Hypothesis: In the mid term, some institutions may use AI to further personalize financial aid offers to optimize yield, potentially increasing the opacity of the offer, acceptance & quota management process.

Far (2046–2051)

In the far term, two rival trajectories emerge for this subsegment of institutions. One path involves a regulatory response to AI-driven opacity, mandating a standardized, machine-readable format for all financial aid offers to ensure transparency. A competing path involves the obsolescence of the "offer letter" as a discrete artifact, replaced by real-time, AI-managed cost-of-attendance dashboards. It remains unknown whether accreditation bodies will view AI-generated personalized pricing as a violation of equitable access standards. The fundamental requirement for a legal offer of admission will likely resist change due to the contractual nature of enrollment.

Hypothesis A: In the far term, regulatory mandates may force a transition to standardized, AI-verifiable financial aid disclosures.

Hypothesis B (competing): In the far term, the traditional offer letter may be replaced by dynamic, AI-managed financial transparency dashboards.

References

1. The Hechinger Report (2026). "How much will that college cost you? Good luck figuring it out." https://hechingerreport.org/how-much-will-that-college-cost-you-good-luck-figuring-it-out-2/

Student Assessment

capability_id: student-assessment · created: 2026-05-10 · revised: 2026-05-10 · refs: 17

Near (2026–2031)

For student assessment broadly, the educational substance is shifting away from a model centered on producing static academic artifacts and answers toward one that prioritizes process, judgment, reflection, and applied thinking [23]. In the process of assessment delivery, there is an emerging pattern of iterative human-AI writing cycles where students use AI as a collaborative partner [1], though this introduces a risk of "cognitive surrender" where students may abandon the logical thinking and research processes the assessment is designed to measure [5]. This shift in delivery mechanism may lead to a homogenization of writing patterns, potentially reducing the ability of assessors to identify individual critical voices during assessment marking and feedback [9]. In assessment administration, the reliability of authorship verification is constrained by the technical failure of AI detection tools, specifically the occurrence of false positives [13] and the existence of "humanizer" tools that bypass detection at a sentence level [15], rendering binary AI detection unreliable [6][12]. For a subset of institutions implementing AI in high-volume or STEM contexts, the process of assessment marking and feedback is incorporating the automation of low-stakes mock exams [8] and the use of confidence-based cascade scoring to manage routine grading [21]. High-stakes, invigilated examinations are unlikely to change in the near term due to the rigid requirements of professional accreditation regimes and the persisting institutional need for authenticated baseline competency [19].

Hypothesis: In the near term, assessment delivery may shift from artifact-production to process-evaluation, while assessment administration is increasingly constrained by the technical failure of authorship detection tools.

Mid (2036–2041)

For student assessment broadly, if current trajectories hold, the educational substance of assessment delivery may shift from the evaluation of a final product to the evaluation of the iterative process of AI interaction and prompting [1][3]. Extrapolating from evidence of cognitive surrender and the failure of detection tools [5][15], the process of marking and feedback may move toward assessing a student's capacity to audit, verify, and correct synthetic errors as the primary proxy for logical reasoning [5][7]. This would transition the capability from measuring product output to measuring the student's ability to steer and validate synthetic agents [14]. For a subset of institutions utilizing multimodal cognitive monitoring, assessment results management may shift toward the longitudinal mapping of cognitive effort and engagement via educational fMRIs and similar multimodal AI [4]. The basic administrative recording of credits and the formal transmission of final results will likely persist unchanged as these are low-variance institutional functions with high regulatory inertia.

Hypothesis: In the mid term, the substance of student assessment may shift from measuring product output to measuring a student's ability to audit and validate synthetic agents.

Far (2046–2051)

For student assessment broadly, the long-term trajectory remains uncertain regarding whether synthetic protocols can fully replace human judgment in certifying higher-order cognitive capabilities [18]. One possible outcome involves the adoption of LLM-based protocols, such as Vantage, to shift assessment delivery toward the continuous, automated measurement of collaboration, creativity, and critical thinking [17]. Conversely, a rival outcome may involve a return to extreme analog, "human-only" assessment environments to combat total cognitive surrender and ensure the social legitimacy of the degree [5]. It remains unknown if AI-generated assessments can be certified as pedagogically valid without human-in-the-loop verification [18]. The fundamental requirement for a third-party institutional certification of competency is likely to resist change, as the social value of the degree depends on an external guarantee of skill.

Hypothesis A: In the far term, assessment delivery may transition to continuous, agent-mediated monitoring of higher-order cognitive processes.

Hypothesis B (competing): In the far term, high-stakes assessment may revert to restricted, non-technological environments to guarantee the authenticity of human cognition.

References

1. The Conversation (2026-03-26). "College students are writing with AI – but a pilot study finds they're not simply letting it write for them." https://theconversation.com/college-students-are-writing-with-ai-but-a-pilot-study-finds-theyre-not-simply-letting-it-write-for-them-276856
2. The Conversation Au (2026-04-01). "Does AI mean more uni students are plagiarising their work?" https://theconversation.com/does-ai-mean-more-uni-students-are-plagiarising-their-work-279565
3. Wonkhe (2026-04-01). "Generative AI exposes longstanding flaws in the use of essay assessments in a mass system." https://wonkhe.com/blogs/generative-ai-exposes-longstanding-flaws-in-the-use-of-essay-assessments-in-a-mass-system/
4. Getting Smart (2026-04-01). "Educational fMRIs: Dynamic Pedagogy and Pedagogical Analysis in the Multimodal AI Era." https://www.gettingsmart.com/2026/04/01/educational-fmris-dynamic-pedagogy-and-pedagogical-analysis-in-the-multimodal-ai-era/
5. Ars Technica (2026-04-03). ""Cognitive surrender" leads AI users to abandon logical thinking, research finds." https://arstechnica.com/ai/2026/04/research-finds-ai-users-scarily-willing-to-surrender-their-cognition-to-llms/
6. Dev To Ai (2026-04-06). "Will AI-Written Papers Be Detected? The Truth About AIGC Detection in 2026." https://dev.to/yu_z_ed6c1c576593c5e7d41f/will-ai-written-papers-be-detected-the-truth-about-aigc-detection-in-2026-5b50
7. Dev To Ai (2026-04-07). "Where Does Truth Live in AI-Generated Code?" https://dev.to/wintrover/where-does-truth-live-in-ai-generated-code-5hmc
8. BBC News (2026-04-07). "'Our teachers use AI to mark mock exams'." https://www.bbc.com/news/articles/c6207xenj1do?at_medium=RSS&at_campaign=rss
9. Hacker News (2026-04-07). "AI may be making us think and write more alike." https://dornsife.usc.edu/news/stories/ai-may-be-making-us-think-and-write-more-alike/
10. Plos One (2026-04-07). "Correction: Evaluating cognitive depth of AI-generated multiple-choice questions with Bloom's Taxonomy." https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0346800
11. Arxiv Cs Cy (2026-04-08). "A Multi-Agent Approach to Validate and Refine LLM-Generated Personalized Math Problems." https://arxiv.org/abs/2604.05160
12. Dev To Ai (2026-04-08). "Your article is well-written. But is it yours?" https://dev.to/dzhuneyt/your-article-is-well-written-but-is-it-yours-59kh
13. Dev To Ai (2026-04-09). "Why AI Detectors Produce False Positives: A Technical Analysis." https://dev.to/laakash/why-ai-detectors-produce-false-positives-a-technical-analysis-2gpc
14. Atlantic (2026-04-10). "This Is What Fully Automated School Looks Like." https://www.theatlantic.com/technology/2026/04/ai-agents-school-education/686754/?utm_source=feed
15. Dev To Ai (2026-04-11). "I Built a Free AI Detector + Humanizer with Sentence-Level Highlighting." https://dev.to/fafa_ai/i-built-a-free-ai-detector-humanizer-with-sentence-level-highlighting-51jj
16. Getting Smart (2026-04-13). "A Call to Action for AI to Promote Mathematical Reasoning." https://www.gettingsmart.com/2026/04/13/a-call-to-action-for-ai-to-promote-mathematical-reasoning/
17. Marktechpost (2026-04-14). "Google AI Research Proposes Vantage: An LLM-Based Protocol for Measuring Collaboration, Creativity, and Critical Thinking." https://www.marktechpost.com/2026/04/13/google-ai-research-proposes-vantage-an-llm-based-protocol-for-measuring-collaboration-creativity-and-critical-thinking/
18. Arxiv Cs Cy (2026-04-14). "Explainability and Certification of AI-Generated Educational Assessments." https://arxiv.org/abs/2604.09622

19. University Affairs (2026-04-14). “Reconsidering the value of exams in the age of AI.” https://universityaffairs.ca/career-advice/reconsidering-the-value-of-exams-in-the-age-of-ai/
20. Arxiv Cs Ai (2026-04-23). “AI to Learn 2.0: A Deliverable-Oriented Governance Framework and Maturity Rubric for Opaque AI in Learning-Intensive Domains.” https://arxiv.org/abs/2604.19751
21. Arxiv Cs Cy (2026-04-23). “Do Small Language Models Know When They’re Wrong? Confidence-Based Cascade Scoring for Educational Assessment.” https://arxiv.org/abs/2604.19781
22. Elearning Industry (2026-04-22). “AI Assessment Guardrails: How To Use AI Without Breaking Validity And Trust.” https://elearningindustry.com/ai-assessment-guardrails-how-to-use-ai-without-breaking-validity-and-trust
23. EDUCAUSE Review (2026-04-23). “Teaching and Learning with AI: Three Priorities.” https://er.educause.edu/podcasts/educause-exchange/teaching-and-learning-with-ai-three-priorities

Student Attraction & Recruitment

capability_id: student-attraction-recruitment · created: 2026-04-25 · revised: 2026-04-25 · refs: 5

Near (2026–2031)

For a subset of institutions adopting AI-driven recruitment tools and those in regions affected by enrollment shifts, the process of prospective student engagement may shift as discovery moves from keyword search to AI-mediated search platforms, necessitating a transition from traditional SEO to AI search optimization [1]. For institutions attending recruitment fairs, the process of prospective student engagement may incorporate AI to convert raw lead data, such as notes and scanned badges, into qualified narratives of student intent and urgency [4]. Simultaneously, the production cost of personalized visual content for domestic and international student recruitment may decrease through the adoption of AI-driven video generation tools [2]. In regions experiencing the long-term effects of the shift to remote learning, the process of domestic student recruitment is constrained by a systemic drain on college enrollment [3]. Within the process of scholarship and bursary management, evidence from some institutions indicates that financial aid offer letters often lack transparency and use confusing terms, creating a significant friction point in prospective student engagement [5]. These changes primarily affect the delivery mechanism of recruitment—how leads are qualified and how financial information is communicated—rather than the educational substance of the offerings. The internal financial governance and audit logic used to allocate scholarship funds is unlikely to change in this period due to rigid regulatory constraints.

Hypothesis: In the near term, a subset of institutions may adjust prospective student engagement to prioritize AI search optimization and automated lead qualification, while communication failures in scholarship management and regional enrollment drains may constrain overall recruitment effectiveness.

Mid (2036–2041)

For the same subset of institutions using AI-driven engagement, if current trajectories hold, the process of prospective student engagement could shift toward the maintenance of structured data schemas designed for machine readability rather than human-centric web narratives [1]. Extrapolating from current lead-intelligence deployments, the conversion of recruitment fair data into qualified narratives may move from an asynchronous post-event task to a real-time automated pipeline [4]. The proliferation of low-cost, AI-generated video may lead to a saturation of personalized outreach, potentially shifting the focus of international student recruitment toward the verification of institutional authenticity [2]. If regional enrollment drains persist, domestic student recruitment may be forced to pivot toward alternative lead-generation mechanisms that bypass traditional pipeline markers [3]. To resolve existing communication failures in scholarship and bursary management, institutions may deploy AI-mediated transparency tools to translate complex financial aid terms into personalized, clear projections for families [5]. However, the institutional requirement for faculty-led program descriptions will likely persist to maintain academic legitimacy.

Hypothesis: In the mid term, prospective student engagement for some institutions may shift toward the production of machine-optimized entity data, while AI may be used to reduce information asymmetry in scholarship and bursary management.

Far (2046–2051)

For institutions relying on AI-mediated prospective student engagement, it remains unknown whether AI search platforms will maintain neutral recommendation architectures or implement proprietary filtering that restricts institutional visibility. One plausible outcome is that the discovery process becomes entirely opaque, where autonomous AI agents match students to institutions based on internal data profiles, rendering public-facing attraction content obsolete. Alternatively, a counterproductive trend may emerge where the saturation of AI-optimized content leads prospective students to distrust digital signals entirely, increasing the relative value of high-touch, human-led recruitment processes. The tension between automated communication and the need for financial transparency in scholarship and bursary management

may result in a regulatory mandate for "human-verified" financial offers to prevent systemic consumer confusion [5]. Despite these shifts, the core legal requirement for institutions to provide an official, binding offer of admission remains a point of resistance against full automation.

Hypothesis A: In the far term, the discovery process for a subset of institutions may be fully mediated by autonomous AI agents, removing the institution's direct control over the attraction process.

Hypothesis B (competing): In the far term, a reaction against AI-saturated recruitment may restore a premium to human-led engagement and verified human signals in the attraction process.

References

1. Dev To Ai (2026-04-06). "What is the difference between AI content and AI search optimised content?" https://dev.to/rank4ai/what-is-the-difference-between-ai-content-and-ai-search-optimised-content-37bc
2. Google AI Blog (2026-04-02). "Create, edit and share videos at no cost in Google Vids" https://blog.google/products-and-platforms/products/workspace/google-vids-updates-lyria-veo/
3. Higher Ed Dive (2026-04-10). "COVID remote learning put drain on college enrollment" https://www.highereddive.com/news/covid-remote-learning-put-drain-on-college-enrollment/816899/
4. Dev To Ai (2026-04-12). "From Notes to Narrative: How AI Automates Trade Show Lead Intelligence" https://dev.to/ken_deng_ai/from-notes-to-narrative-how-ai-automates-trade-show-lead-intelligence-1239
5. The Hechinger Report (2026-04-20). "How much will that college cost you? Good luck figuring it out" https://hechingerreport.org/how-much-will-that-college-cost-you-good-luck-figuring-it-out-2/

Student Enrolment

capability_id: student-enrolment · created: 2026-04-10 · revised: 2026-04-10 · refs: 1

Near (2026–2031)

In the United States, programme enrolment volumes for first-year students have decreased [1]. This decline is linked to a lower rate of completion for financial aid applications (FAFSA) and standardized testing (ACT) [1]. These frictions at the entry point of the programme enrolment process reduce the pool of eligible candidates entering the pipeline [1]. Because this drain occurs at the initial sign-up phase, module selection management is unlikely to change in the near term as it is a downstream process. Student induction management remains largely unaffected by these volume shifts in the immediate term [1].

Hypothesis: In the near term, programme enrolment in the US region may experience continued volatility as institutions respond to changes in student application behaviors.

Mid (2036–2041)

In the United States, if current trajectories of declining first-year sign-ups hold, institutions may alter the delivery mechanism of programme enrolment to reduce friction [1]. This could involve the deployment of automated tools to assist students in navigating FAFSA and ACT requirements, shifting the process from a passive administrative task to an active, supported workflow. Extrapolating from the volume drain, institutions may also modify student induction management to prioritize retention of a smaller entering cohort [1]. The fundamental regulatory requirements for financial aid eligibility will likely persist unchanged due to government oversight.

Hypothesis: In the mid term, institutions in the affected region may adopt more automated delivery mechanisms for programme enrolment to mitigate volume loss.

Far (2046–2051)

In the United States, the long-term elasticity of student demand relative to remote learning formats remains unknown [1]. If the drain on traditional programme enrolment persists, the capability may diverge into two rival outcomes. One outcome involves the complete automation of the enrolment pipeline, where institutional barriers are minimized to maximize throughput. A competing outcome involves a return to high-touch, human-centric enrolment processes to restore social legitimacy and perceived value [1]. Legal requirements for institutional accreditation will resist change even at this horizon, as they provide the basis for enrolment legitimacy.

Hypothesis A: In the far term, programme enrolment may transition to a continuous, fully automated delivery mechanism to minimize student friction.

Hypothesis B (competing): In the far term, institutions may shift toward high-touch, human-led enrolment processes to counteract the perceived depersonalization of remote learning.

References

1. Higher Ed Dive (2026). “COVID remote learning put drain on college enrollment.” https://www.highereddive.com/news/covid-remote-learning-put-drain-on-college-enrollment/816899/

Student Support & Wellbeing Management

capability_id: student-support-wellbeing-management · created: 2026-05-02 · revised: 2026-05-02 · refs: 9

Near (2026–2031)

For a subset of institutions and specific student populations incorporating AI-mediated workflows, shifts are emerging in both the delivery and substance of support processes. In the process of personal learning management, the delivery mechanism is shifting toward AI-automated synthesis of study materials through the deployment of tools such as Adobe Acrobat Spaces [4]. Within academic skill development, the educational substance is shifting toward the regulation of "cognitive partnership," as students move from simple generation toward iterative brainstorming [1] and the adoption of frameworks like TACO to prevent AI from substituting for logical thinking [8]. This shift is a response to the risk of "cognitive surrender," where users abandon independent reasoning [2] and exhibit linguistic homogenization in their thought and writing patterns [5]. In the process of disability support management, delivery is shifting toward the use of AI for drafting Individualized Education Programs (IEPs) and coursework support [9], though these deployments currently lack systematic safeguarding and policy frameworks [9]. Within student health & wellbeing, delivery is incorporating AI despite workforce tension between clinical enthusiasm and fear of displacement [3], while the substance of triage is adapting to address new psychological risks, including AI-induced delusions [6]. In the region of Florida, this process is being formally constrained by state-directed partnerships to shield families from AI harm [7]. The processes of student financial advice and housing advice are unlikely to change in this horizon because they remain tied to external regulatory frameworks and static administrative constraints.

Hypothesis: In the near term, academic skill development for a subset of students may transition toward the formal regulation of cognitive partnership to mitigate cognitive surrender, while disability support management may see increased AI-assisted IEP production alongside a critical lag in safeguarding policies.

Mid (2036–2041)

For the same subset of institutions and regions, if current trajectories hold, the academic skill development process may formalize the auditing of machine-generated logical structures as a core competency, extrapolating from early frameworks that distinguish between cognitive support and cognitive substitution [8]. Within disability support management, the implementation of "Responsible Inference Engines" may evolve from a policy recommendation into a standard operational requirement for safeguarding students with learning differences [9]. In student health & wellbeing, AI-driven documentation and triage may further integrate into the workforce, provided that institutional protocols resolve the tension between professional fear and operational efficiency [3]. In regions implementing standardized harm reporting, such as Florida, these mechanisms may evolve into permanent regulatory requirements for student wellbeing management [7]. The human-led provision of personal tutors is likely to persist unchanged as the primary mechanism for emotional scaffolding and motivational support, as these require a level of social legitimacy and human empathy not evidenced in current AI deployments.

Hypothesis: In the mid term, institutions supporting AI-integrated workflows may establish formalized standards for the regulation of cognitive partnership and the safe deployment of inference engines for disability support.

Far (2046–2051)

For the subsegments of support focused on literacy, disability, and mental health, long-term outcomes depend on whether AI integration enhances or erodes foundational cognitive and therapeutic bonds. One plausible outcome is that academic skill development evolves to prioritize high-level conceptual orchestration, treating prose as a utility, while health services move toward high-efficiency, AI-managed triage and recording [3][4]. A competing outcome is that the risk of cognitive surrender [2] leads to a systemic atrophy of independent reasoning, necessitating a return to strictly human-mediated academic skill development to

restore cognitive autonomy. It remains unknown whether synthetic psychological support can ever achieve the social legitimacy required for high-stakes mental health interventions. The requirement for human accreditation and signatures in disability support and high-stakes health certifications is likely to resist change even at this horizon due to legal and professional liability constraints.

Hypothesis A: In the far term, support capabilities for a subset of students may shift toward high-efficiency conceptual orchestration and AI-managed triage.

Hypothesis B (competing): In the far term, the prevalence of cognitive surrender may force a structural return to human-only mediation in academic skill development to preserve foundational reasoning.

References

1. The Conversation (2026-03-26). "College students are writing with AI – but a pilot study finds they're not simply letting it write for them." https://theconversation.com/college-students-are-writing-with-ai-but-a-pilot-study-finds-theyre-not-simply-letting-it-write-for-them-276856
2. Ars Technica (2026-04-03). ""Cognitive surrender" leads AI users to abandon logical thinking, research finds." https://arstechnica.com/ai/2026/04/research-finds-ai-users-scarily-willing-to-surrender-their-cognition-to-llms/
3. NPR (2026-04-07). "AI in the mental health care workforce is met with fear, pushback — and enthusiasm." https://www.npr.org/2026/04/07/nx-s1-5771707/mental-health-care-workforce-artificial-intelligence-ai
4. TechCrunch (2026-04-07). "Adobe launches Acrobat Spaces, a free AI-powered study tool for students." https://techcrunch.com/2026/04/07/adobe-launches-acrobat-spaces-a-free-ai-powered-study-tool-for-students/
5. Hacker News (2026-04-07). "AI may be making us think and write more alike." https://dornsife.usc.edu/news/stories/ai-may-be-making-us-think-and-write-more-alike/
6. Daily Ai (2025-05-09). "ChatGPT Is Making People Think They're Gods and Their Families Are Terrified." https://dailyai.com/2025/05/chatgpt-is-making-people-think-theyre-gods-and-their-families-are-terrified/
7. Future Of Life (2026-03-09). "Governor DeSantis Directs Florida State Agencies to Partner with Future of Life Institute to Shield Families from AI Harm." https://futureoflife.org/press-release/desantis-directs-florida-agencies-to-partner-with-fli/
8. Arxiv Cs Cy (2026-04-22). "Students Know AI Should Not Replace Thinking, but How Do They Regulate It? The TACO Framework for Human-AI Cognitive Partnership." https://arxiv.org/abs/2604.18737
9. Getting Smart (2026-04-22). "Responsible Inference Engines: Safeguarding Students with Learning Differences in the AI Era." https://www.gettingsmart.com/2026/04/22/responsible-inference-engines-safeguarding-students-with-learning-differences-in-the-ai-era/

Supporting Services

capability_id: supporting-services · created: 2026-05-11 · revised: 2026-05-11 · refs: 3

Near (2026–2031)

For specific subsegments of supporting services, primarily food and retail management and venue management, the delivery mechanisms for logistics and monitoring are shifting toward automated digital systems. In food and retail management, the implementation of automated approval systems for supply chain logistics creates a dependency where physical goods cannot move without digital validation, which may lead to the physical waste of perishables when systems fail to approve shipments [1]. In venue management, the integration of AI video analytics transforms traditional CCTV into intelligent monitoring systems capable of detecting real-time risks through object detection and tracking [3]. Furthermore, the application of predictive intelligence allows these venues to shift from reactive observation toward forecasting movement patterns and density changes [2]. Processes such as religious support and child care management are unlikely to adopt these specific automated frameworks in the near term because their institutional value is derived from human presence and social legitimacy.

Hypothesis: In the near term, food and retail management may experience increased operational fragility due to digital validation failures, while venue management may shift toward proactive risk detection and predictive spatial monitoring.

Mid (2036–2041)

For the food, retail, and venue management subsegments, if current trajectories of digital integration hold, the stability of these services will depend on the resilience of the underlying software. Extrapolating from current failures in automated approvals, the risk of systemic waste in food and retail management could become a structural vulnerability unless institutions implement manual override protocols [1]. In venue management, predictive intelligence and video analytics may evolve from monitoring density to automating real-time resource allocation and crowd control mechanisms [2][3]. However, the educational substance of complaint and compliment management is likely to remain unchanged, as the resolution of human grievances typically requires social empathy and discretionary authority.

Hypothesis: In the mid term, food logistics may face persistent risks of "digital waste," while venue management may integrate predictive analytics into automated operational resource allocation.

Far (2046–2051)

For the food, retail, and venue management subsegments, the long-term outcome depends on whether the industry optimizes for centralized efficiency or distributed resilience. One possibility involves the development of high-resilience, autonomous validation systems that mitigate the single points of failure currently causing supply chain waste [1]. A rival outcome is a state of chronic operational instability where the gap between digital approval and physical reality creates permanent inefficiencies in food and retail management. It remains unknown if the cumulative cost of system-induced waste will force a return to non-digital, localized procurement. Processes like religious support will likely resist this trajectory throughout this horizon as their primary function is the provision of human-centric spiritual care.

Hypothesis A: In the far term, these subsegments may achieve autonomous, high-resilience operational states via distributed intelligence and predictive validation.

Hypothesis B (competing): In the far term, supporting services may experience structural degradation and chronic waste due to an irreversible dependence on rigid, centralized digital approval mechanisms.

References

1. Science Daily Comp (2026-04-05). "Truckloads of food are being wasted because computers won't approve them." https://www.sciencedaily.com/releases/2026/04/260403224505.htm
2. Dev To Ai (2026-04-22). "Architecting Predictive Intelligence for Smart Venues." https://dev.to/t

anisha_khiara_4b390b7946/architecting-predictive-intelligence-for-smart-venues-1lkj
3. Dev To Ai (2026-04-23). “Integrating AI Video Analytics into Venue Operations.” https://dev.to/t anisha_khiara_4b390b7946/integrating-ai-video-analytics-into-venue-operations-2m09

Teaching & Learning Delivery

capability_id: teaching-learning-delivery · created: 2026-05-13 · revised: 2026-05-13 · refs: 18

Near (2026–2031)

For teaching & learning delivery broadly, the educational substance is shifting toward skills-based learning frameworks to address rapid workplace volatility [13] and a prioritization of learning processes over the production of academic artifacts [19]. Within teaching & learning content management, the delivery mechanism is shifting from manual authoring to automated multimodal generation, including instructional videos [6], presentations derived from URLs [16], automated multimodal illustrations [7], and localized materials for Career and Technical Education (CTE) [1]. This process is further modified by the transition to agentic co-authoring tools that directly modify documents [20] and the use of expressive synthetic speech to reduce reliance on human voice-overs in asynchronous materials [21]. In non presence based teaching, delivery mechanisms are being modified by AI avatars that convert static documents into interactive experiences [14], personalized practice systems that maintain learning rate regularity at scale [3], and specialized AI tutoring for coding [2] and music exam preparation [4]. For institutions implementing structured AI integration in software engineering and STEM, presence based teaching is being modified by frameworks that establish a "cognitive partnership" to ensure AI does not replace the student's internal thinking process [15][17]. The core social mediation and professional socialization functions of presence based teaching are unlikely to change yet because the social legitimacy of the human instructor remains a primary requirement for institutional accreditation.

Hypothesis: In the near term, the capability shifts toward a process-oriented educational substance supported by a delivery mechanism of automated multimodal content and adaptive non presence based interfaces.

Mid (2036–2041)

For institutions adopting automated delivery models and STEM-focused subsegments, the delivery mechanism of non presence based teaching may shift toward agent-led orchestration of student pathways if current trajectories in AI tutoring and autonomous schooling hold [2][10]. Within teaching & learning content management, the process may evolve from a creation-centric activity into a verification-heavy auditing role, provided that governance frameworks for quality, accuracy, and trust are standardized [5][11]. If the risk of "cognitive bypass"—where students use AI as a substitute for critical thinking—persists, presence based teaching could shift toward constrained delivery models that intentionally limit AI autonomy to protect the cognitive substance of the learning process [17]. The requirement for human-led mentorship in high-stakes professional socialization is likely to persist because it is embedded in the regulatory regimes of professional licensing and accreditation.

Hypothesis: In the mid term, content management may transition into a verification and governance function, while non presence based teaching may shift toward agent-orchestrated pathways.

Far (2046–2051)

For the same subset of institutions, the delivery of teaching and learning may diverge based on how the tension between efficiency and cognitive development is resolved. One outcome involves the full transition of non presence based teaching to autonomous agent-led systems for knowledge acquisition, leaving human instructors to focus exclusively on high-level synthesis and professional identity [10]. A competing outcome involves a strategic re-centering on high-touch, human-only presence based teaching as a prestige marker in response to the ubiquity of automated delivery. It remains unknown whether the long-term use of AI partners will erode or enhance the human capacity for independent mathematical and logical reasoning [12]. The institutional requirement for human-validated credentialing resists change even at this horizon due to the necessity of social trust in professional certification.

Hypothesis A: In the far term, the capability bifurcates into fully autonomous knowledge acquisition and high-prestige, human-only mentorship.

Hypothesis B (competing): In the far term, the capability stabilizes around a hybrid model where AI agents handle all delivery mechanisms while humans retain sole authority over the educational substance.

References

1. Elearning Industry (2026-04-05). "Leveraging AI And Automation To Streamline CTE Content Creation And Localization." https://elearningindustry.com/leveraging-ai-and-automation-to-streamline-cte-content-creation-and-localization
2. The Hechinger Report (2026-04-06). "The quest to build a better AI tutor." https://hechingerreport.org/proof-points-ai-tutor-python/
3. Arxiv Cs Cy (2026-04-07). "Personalized AI Practice Replicates Learning Rate Regularity at Scale." https://arxiv.org/abs/2604.03246
4. Dev To Ai (2026-04-07). "Automating Excellence: AI for Custom Music Exam & Recital Prep." https://dev.to/ken_deng_ai/automating-excellence-ai-for-custom-music-exam-recital-prep-2e2d
5. Elearning Industry (2026-04-07). "A 4-Step Framework For Using AI Transparently In Educational Content." https://elearningindustry.com/a-4-step-framework-for-using-ai-transparently-in-educational-content
6. Google AI Blog (2026-04-02). "Create, edit and share videos at no cost in Google Vids." https://blog.google/products-and-platforms/products/google-vids-updates-lyria-veo/
7. Arxiv Cs Cy (2026-04-08). "EduIllustrate: Towards Scalable Automated Generation Of Multimodal Educational Content." https://arxiv.org/abs/2604.05005
8. Arxiv Cs Cy (2026-04-08). "A Multi-Agent Approach to Validate and Refine LLM-Generated Personalized Math Problems." https://arxiv.org/abs/2604.05160
9. EDUCAUSE Review (2026-04-08). "Building AI Initiatives with Bottom-Up Champions." https://er.educause.edu/articles/2026/4/building-ai-initiatives-with-bottom-up-champions
10. Atlantic (2026-04-10). "This Is What Fully Automated School Looks Like." https://www.theatlantic.com/technology/2026/04/ai-agents-school-education/686754/?utm_source=feed
11. Elearning Industry (2026-04-10). "AI-Generated Learning Content Governance: Ensuring Quality, Accuracy, And Trust." https://elearningindustry.com/ai-generated-learning-content-governance-ensuring-quality-accuracy-and-trust
12. Getting Smart (2026-04-13). "A Call to Action for AI to Promote Mathematical Reasoning." https://www.gettingsmart.com/2026/04/13/a-call-to-action-for-ai-to-promote-mathematical-reasoning/
13. Elearning Industry (2026-04-15). "The Case For Skills-Based Learning In A Rapidly Changing Workplace." https://elearningindustry.com/the-case-for-skills-based-learning-in-a-rapidly-changing-workplace
14. Elearning Industry (2026-04-21). "How AI Avatars Are Turning Static Documents Into Interactive Learning Experiences." https://elearningindustry.com/how-ai-avatars-are-turning-static-documents-into-interactive-learning-experiences
15. Arxiv Cs Cy (2026-04-21). "Bringing AI into the Classroom: A Structured Approach for Integrating AI into Software Engineering Education." https://arxiv.org/abs/2604.16328
16. Dev To Ai (2026-04-21). "Create Presentations from Any URL Automatically Using AI." https://dev.to/ai_presentation_d9935ff95/create-presentations-from-any-url-automatically-using-ai-5hho
17. Arxiv Cs Cy (2026-04-22). "Students Know AI Should Not Replace Thinking, but How Do They Regulate It? The TACO Framework for Human-AI Cognitive Partnership." https://arxiv.org/abs/2604.18737
18. Getting Smart (2026-04-22). "Responsible Inference Engines: Safeguarding Students with Learning Differences in the AI Era." https://www.gettingsmart.com/2026/04/22/responsible-inference-engines-safeguarding-students-with-learning-differences-in-the-ai-era/
19. EDUCAUSE Review (2026-04-23). "Teaching and Learning with AI: Three Priorities." https://er.educause.edu/podcasts/educause-exchange/teaching-and-learning-with-ai-three-priorities
20. The Register (2026-04-23). "Microsoft gives your Word documents an AI co-author you didn't ask for." https://go.theregister.com/feed/www.theregister.com/2026/04/23/microsoft_gives_your_wo

rd_documents/
21. Google DeepMind (2026-04-15). "Gemini 3.1 Flash TTS: the next generation of expressive AI speech." https://deepmind.google/blog/gemini-3-1-flash-tts-the-next-generation-of-expressive-ai-speech/